Log-$\mathcal{S}$-unit Lattices Using Explicit Stickelberger Generators to Solve Approx Ideal-SVP

Olivier Bernard [1,2], Andrea Lesavourey [1], Tuong-Huy Nguyen[1,3], and Adeline Roux-Langlois [1]

[1] Univ Rennes, CNRS, IRISA, France
olivier.bernard@normalesup.org,
{andrea.lesavourey, tuong-huy.nguyen, adeline.roux-langlois}@irisa.fr
[2] Thales, Gennevilliers, France
[3] DGA Maîtrise de l'Information, Bruz, France

Abstract. In 2020, Bernard and Roux-Langlois introduced the Twisted-PHS algorithm to solve Approx-Svp for ideal lattices on any number field, based on the PHS algorithm by Pellet-Mary, Hanrot and Stehlé. They performed experiments for prime conductors cyclotomic fields of degrees at most 70, one of the main bottlenecks being the computation of a log-$\mathcal{S}$-unit lattice which requires subexponential time.
Our main contribution is to extend these experiments to cyclotomic fields of degree up to 210 for most conductors m. Building upon new results from Bernard and Kučera on the Stickelberger ideal, we use explicit generators to construct full-rank log-$\mathcal{S}$-unit sublattices fulfilling the role of approximating the full Twisted-PHS lattice. In our best approximate regime, our results show that the Twisted-PHS algorithm outperforms, over our experimental range, the CDW algorithm by Cramer, Ducas and Wesolowski, and sometimes beats its asymptotic volumetric lower bound. Additionally, we use these explicit Stickelberger generators to remove almost all quantum steps in the CDW algorithm, under the mild restriction that the plus part of the class number verifies $h_m^+ \leq O(\sqrt{m})$.

Keywords: Ideal lattices, Approx-SVP, Stickelberger ideal, S-unit attacks, Twisted-PHS algorithm

1 Introduction

The ongoing NIST Post-Quantum Cryptography competition illustrates the importance of the *Learning With Errors* (Lwe) problem as an intermediate building block for a wide variety of cryptographic schemes. Most of these cryptographic schemes rely on a structured version of the Lwe problem allowing for much more satisfactory performance, compared to schemes based on the unstructured Lwe problem. The first structured variant of Lwe, later known as the Ring-Lwe problem, is shown to be at least as hard as the *Approximate Shortest Vector Problem* on ideal lattices (Approx-id-Svp) using quantum worst-case to average-case reductions [SSTX09,LPR10]. One important matter is to determine whether

using this structured version of Lwe could lower the hardness hypothesis of the scheme. Notably, an efficient solver for Approx-id-SVP would render the worst-case to average-case reduction to Ring-Lwe meaningless as a security argument. Note however that even in this case, this would not directly imply an efficient solver for the Ring-Lwe problem.

In the case of arbitrary lattices, Approx-Svp is a well-studied hard problem. It consists in finding relatively short vectors of a given lattice, within an approximation factor of the shortest vector. The best theoretical trade-off between runtime and approximation factor is known as Schnorr's hierarchy [Sch87]: one can reach, for any $\alpha \in (0,1)$, an approximation factor $2^{\widetilde{O}(n^{\alpha})}$ in time $2^{\widetilde{O}(n^{1-\alpha})}$. The closest known practical algorithm to this trade-off is the BKZ algorithm [SE94], a generalization of the well-known LLL algorithm [LLL82]. In the particular case of ideal lattices, i.e., lattices that correspond to ideals of the ring of integers $\mathcal{O}_K$ of a number field K, one could hope that the best reduction algorithms would remain those associated with arbitrary lattices. However, this simplifying assumption seems questionable, since the underlying number-theoretic structure is precisely what makes Ring-Lwe a more efficient building block. Thus, going beyond the BKZ algorithm and estimating the hardness of Approx-id-Svp using algebraic ideas has gathered more attention, starting by works from [EHKS14,CGS14,BS16,CDPR16]. Earlier works aimed at the more restricted case of Approx-id-Svp for principal ideals. A strategy for this case is devised as a two parts algorithm. The first part requires solving the Principal Ideal Problem (Pip), i.e., finding any generator of the ideal; the second part aims at finding the shortest one, by solving a Closest Vector Problem (Cvp) in the so-called *log-unit lattice*. This shortest generator is expected to solve Approx-Svp for a sufficiently small approximation factor. Ultimately, for the particular case of cyclotomic fields of prime power conductors, [CDPR16] proved that Approx-id-Svp on principal ideals is solvable in quantum polynomial time, but only reaching an approximation factor $2^{\widetilde{O}(\sqrt{n})}$.

Subsequent works in a more general case can be divided in two different paths. The first one [CDW17,CDW21] aimed at extending the results from [CDPR16] to arbitrary ideal lattices over any cyclotomic fields, while still reaching in quantum polynomial time an approximation factor $2^{\widetilde{O}(\sqrt{n})}$. One of their contributions is to reduce the arbitrary ideal case to the principal ideal case by solving the *Close Principal Multiple Problem* (Cpmp): given an ideal $\mathfrak{b}$, one computes an ideal $\mathfrak{c}$ of small algebraic norm s.t. $\mathfrak{bc}$ is a principal ideal. In order to ensure that $\mathfrak{c}$ has a small norm, a new key technical ingredient, specific to cyclotomic fields, was the use of the Stickelberger lattice, which has good geometric properties. Then, the results from [CDPR16] are applied to $\mathfrak{bc}$ to obtain a candidate short element of $\mathfrak{b}$, using the fact that $\mathfrak{c}$ has a small norm. The concrete consequences of this method were experimented in [DPW19], under different regimes which mainly differ upon which Cvp solver is used. The first regime (called "Naive") uses Babai's Nearest Plane algorithm, whereas the second regime uses a heuristic Cvp algorithm relatively to *ad hoc* pseudo-norms. From these experiments, the asymptotic performance of those decoding algorithms was estimated, which

led to simulated approximation factors reached by the CDW algorithm. Finally, given experimentally verified constants, a volumetric lower bound was derived for the approximation factors that could be reached in the best scenario. According to this lower bound, the CDW algorithm is expected to beat the BKZ_{300} algorithm for cyclotomic fields of degrees at least larger than 7000. Since NIST submissions based on structured lattices rely on cyclotomic fields of degree at most 1024, this could be perceived as somewhat reassuring.

The second path is explored in [PHS19,BR20]. Those works, applying to arbitrary number fields, replace the two reductions steps from [CDW21] with a single Cvp instance, so as to find a principal multiple ideal which is not only of small algebraic norm, but is also generated by a small element. A key ingredient achieving this is to use a generalization of the units of $\mathcal{O}_K$, called $\mathcal{S}$-units; this formalism was an underlying feature of [PHS19] and was later made explicit in [BR20]. The PHS algorithm splits into a preprocessing phase and a query phase. The preprocessing phase consists in preparing the decoding of a particular lattice depending only on the number field K, *via* the computation of a hint following Laarhoven's Cvp with preprocessing algorithm [Laa16], which takes exponential time. Then, any Approx-id-Svp instance in K can be interpreted as an Approx-Cvp instance in this lattice, efficiently solved thanks to the hint. Up to the preprocessing, the query phase yields new time/quality trade-offs: as in [CDW21] for cyclotomic fields, it reaches approximation factor $2^{\widetilde{O}(\sqrt{n})}$ in quantum polynomial time; however, the PHS algorithm also allows for better trade-offs than Schnorr's hierarchy, from polynomial to $2^{\widetilde{O}(\sqrt{n})}$ approximation factors. On the downside, the computation of the lattice itself takes classically subexponential time, which is a serious obstacle for studying their geometry and obtaining concrete asymptotic estimations as was done in [DPW19] for the CDW algorithm.

Then, [BR20] introduced Twisted-PHS, a "Twisted" version of the PHS algorithm whose main difference lies in a fundamental modification of the underlying lattice, thanks to a natural normalization coming from the *Product Formula*. The problem of finding a short vector is expected to be better encoded within this new lattice, ultimately leading to smaller outputs. Even though the proven trade-offs between runtime and approximation factor remain the same for the Twisted-PHS algorithm as for the PHS algorithm, very significant improvements have been experimentally illustrated in [BR20, Fig. 5.3], showing much better approximation factors compared to the PHS algorithm for number fields of degree up to 60, where Laarhoven's Cvp algorithm is replaced in practice by Babai's Nearest Plane algorithm [Bab86]. These were to our knowledge the first experimental evidence of the geometric peculiarity of normalized log-$\mathcal{S}$-unit lattices and of the practical potential of this type of attack. In this practical version, experiments are solely limited by the classical complexity of computing the lattice.

Unfortunately, the attained dimensions, up to 60, are not sufficient to assess the practical limits of the Twisted-PHS algorithm: its heuristic analysis [BR20] could give only a loose upper bound, or miss unexpected performance in practical dimensions due to its asymptotic nature, even in the cryptographical range.

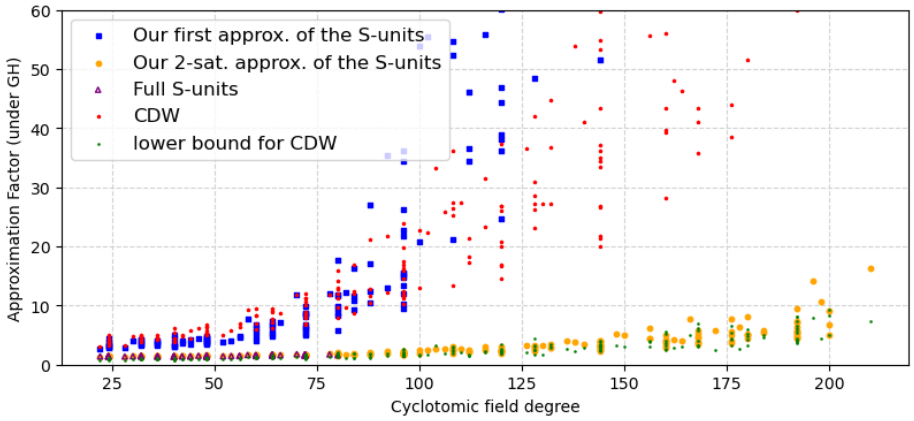

Fig. 1.1 – Average of approximation factors achieved by our implementation of Twisted-PHS, using log-$\mathcal{S}$-unit sublattices in cyclotomic fields over random simulated instances, compared to those achieved by CDW [DPW19], assuming the Gaussian Heuristic throughout all instances.

Our contributions. We develop theoretical and practical improvements regarding algorithms for solving Approx-id-Svp, in both lines of work following the CDW algorithm and the Twisted-PHS algorithm. Even though the hardness of the Approx-id-Svp does not concretely impact the security of cryptographic schemes, it is important to get a better understanding of both approaches, which are the only ones successfully exploiting the structure of a lattice.

Our core ingredient is the introduction of a full-rank family of independent $\mathcal{S}$-units, whose algebraic properties are proven in §3. In §4, we use this family to remove most quantum steps of the CDW algorithm, leaving only one step during a preprocessing phase done once for any given field, and one step for each query.

In §5, this family allows us to achieve experiments on algorithms in the (Twisted-)PHS family, for most cyclotomic fields of dimension up to 210. By comparison, previous experiments [DPW19,BR20] only considered cyclotomic fields of conductors $m = p > 2$ prime and $m = 2^e > 2$. Our work comes with an improved implementation of the initial Twisted-PHS algorithm, allowing us to extend the experiments conducted in [BR20] up to dimension 80 and for all cyclotomic fields. It also includes different regimes of approximation for this algorithm, using sublattices of the log-$\mathcal{S}$-unit lattice obtained thanks to our new construction beyond dimension 80 up to 210. These regimes yield concrete upper bounds for the approximation factors that could be reached by the full Twisted-PHS algorithm up to dimension 210, as illustrated in Fig. 1.1:

1. The depicted approximation factors were estimated using the *Gaussian Heuristic*, matching the *exact ones* obtained by [BR20] without this hypothesis.

2. Our best approximate regime yields approximation factors that are comparable (sometimes even smaller) to the asymptotical volumetric lower bound regime of the CDW algorithm.

In [DPW19], it was already noted that the PHS approach should outperform the lower bound, but at the cost of computing Laarhoven's hint in exponential time. Our work show that for medium dimensions, where asymptotical results should start to be meaningful, the Twisted-PHS algorithm is at least comparable to the CDW lower bound, though without this exponential hint precomputation.

As suggested in [BR20], and illustrated in small dimensions, the Twisted-PHS algorithm performance may be explained by the peculiar geometric nature of the log-$\mathcal{S}$-unit lattice. In our work, this is confirmed by the computations of several geometrical parameters on the basis obtained by our implementation, across all considered cyclotomic fields, sublattices and factor bases. This specificity, observed in a wide variety of regimes and even in medium dimensions, suggest a deeper explanation, a possibility recently explored by Bernstein and Lange [BL21]. We provide a full implementation of all our experiments at `https://github.com/ob3rnard/Tw-Sti`.

Technical overview. In [BR20], the log-$\mathcal{S}$-unit lattice needed for the preprocessing phase was built using generic number theory tools. Our main idea is to shortcut this generic computation by considering a maximal family $\mathfrak{F}$ of independent $\mathcal{S}$-units, where $\mathcal{S}$ verifies some conditions (detailed in §3), leading to sublattices of the log-$\mathcal{S}$-unit lattice. The family $\mathfrak{F}$ is composed of three parts:

1. Circular units, also known as cyclotomic units, e.g. in [Was97, §8];
2. Generators coming from the explicit proof of Stickelberger's theorem proof;
3. Real $\mathcal{S}$-units coming from the maximal real subfield K_m^+ of K_m, where K_m is the cyclotomic field of conductor m.

The first two parts are classically easy to compute. In particular, the effectiveness of the second part comes from two recent results of [BK21]: the knowledge of an explicit *short* $\mathbb{Z}$*-basis* of the Stickelberger ideal for *any* conductor [BK21, Th. 3.6], and the effective computations of generators corresponding to these short relations using Jacobi sums [BK21, §5]. On the contrary, the last part still relies on generic number theory tools which are classically costly, but are now performed in a number field of half degree, which propels us to degree 210.

As an important theoretical contribution, we prove in Th. 3.11 that $\mathfrak{F}$ is indeed a full-rank family of multiplicatively independent $\mathcal{S}$-units, by computing explicitly its (finite) index in the full $\mathcal{S}$-unit group. This can be seen as a generalization of the strategy of [CDW17, Def. 2] to obtain a full-rank lattice of class relations, restricted to the relative class group. In particular, our result proves the experimentally conjectured value [DPW19, Rem. 3] of the index of their family.

Finally, the index of $\mathfrak{F}$ contains a large power of 2 that can be removed using classical 2-saturation techniques of §3.5, leading to a family $\mathfrak{F}_{\mathsf{sat}}$. We then use the explicit knowledge of these special $\mathcal{S}$-units in two different situations.

Theoretical improvements of the CDW algorithm. In §4, we remove almost all quantum steps of the CDW algorithm while still guaranteeing its approximation

factor [CDW21, Th. 5.1], at the small price of restricting to cyclotomic fields s.t. $h_m^+ \leq O(\sqrt{m})$ ([BLNR21, Hyp. B.1]), where h^+ denotes the plus part of the class number (defined in §2.2), whereas [CDW21, Ass. 2] uses $h_m^+ \leq \text{poly}(m)$.

For that purpose, we state an equivalent rewriting of [CDW21, Alg. 7], making explicit some hidden steps useful for subsequent modifications. Then, the explicit Stickelberger generators and real $\mathcal{S}$-units are used to remove the last call to the quantum Pip solver. Finally, considering the module of *all* real class group relations allows us to remove the quantum random walk mapping any ideal of K_m into the relative class group. This last part uses our Th. 3.11 and needs [BLNR21, Hyp. B.1] to obtain the same bound on the approximation factor.

Only two quantum steps remain: the first is performed once to compute real $\mathcal{S}$-units in K_m^+, of degree only half, the second is solving the ClDL for each query.

Experimenting the Twisted-PHS algorithm in medium dimensions. We apply Twisted-PHS [BR20] on our full-rank sublattices of the log-$\mathcal{S}$-unit lattice, yielding *approximated* regimes of the Twisted-PHS algorithm. Up to degree 210, for most conductors, the newly implemented algorithm is used to compute the sublattices associated with $\mathfrak{F}$ and $\mathfrak{F}_{\mathsf{sat}}$, for varying subsets $\mathcal{S}$ according to the number of Galois orbits of totally split primes used. In particular, we explicitly compute the Stickelberger generators and real generators of $\mathfrak{F}$ and effectively perform the 2-saturation of $\mathfrak{F}$ to get $\mathfrak{F}_{\mathsf{sat}}$. Up to degree 80, the whole log-$\mathcal{S}$-unit lattice is also computed, corresponding to a fundamental system $\mathfrak{F}_{\mathsf{su}}$ of $\mathcal{S}$-units. This last computation of $\mathfrak{F}_{\mathsf{su}}$ remains unfeasible at higher dimensions. We evaluate the geometry of all these lattices with standard indicators described in §2.5: the root-Hermite factor δ_0, the orthogonality defect δ and the logarithm of the Gram-Schmidt norms. We consistently observe the same phenomena already pointed out in [BR20, §5.1 and 5.2], that indicate close to orthogonal lattices.

Next, since computing ClDL solutions for random ideals quickly becomes intractable, we simulate this step by sampling random outputs similarly to what was done in [DPW19, Hyp. 8]. Given those targets and the preprocessed lattices associated with $\mathfrak{F}$, $\mathfrak{F}_{\mathsf{sat}}$ and $\mathfrak{F}_{\mathsf{su}}$, we evaluate the approximation factors reached by these different regimes, by assuming the Gaussian Heuristic. These two assumptions, i.e., using simulated targets and the Gaussian Heuristic, are validated by the fact that up to degree 80, where it is feasible to compute the full $\mathcal{S}$-unit group generated by $\mathfrak{F}_{\mathsf{su}}$, our approximation factors match the *exact* approximation factors obtained in [BR20, Fig. 1.1], where those heuristics were not used. Finally, we compare our results to the approximation factors obtained by the CDW algorithm [CDW21] in the "Naive" regime of [DPW19], under the same working assumptions as above. We observe that in our best approximate regime, using $\mathfrak{F}_{\mathsf{sat}}$, our estimated approximation factors are close, and sometimes smaller, than the theoretical lower bound derived in [DPW19]. This suggests that the crossover with BKZ_{300} could be lower than expected for the Twisted-PHS algorithm.

Relations to other works related to $\mathcal{S}$-units. Some recent mathematical results regarding the Stickelberger lattice were established in [BK21]. The authors

described, for *any* conductor, an easily computable *short basis* for this lattice, and how to explicitly compute the associated principal ideal generators through Jacobi sums. In our work, this result is brought into fruition to solve Approx-id-SVP. The completion of this short basis into a full-rank lattice of class relations, the effective computation of the explicit generators and the 2-saturation of these elements, yielded the different approximated regimes of Twisted-PHS and allowed us to remove many quantum steps from the CDW algorithm.

In a talk on August 2021 at SIAM Conference,[4] Bernstein announced a joint work with Eisenträger, Rubin, Silverberg and van Vredendaal, by illustrating the construction of small $\mathcal{S}$-units using Jacobi sums that lead to an "$\mathcal{S}$-unit attack" in the power-of-2 conductor case up to degree 64, assuming $h_{2^e}^+ = 1$. The talk also announced a paper that has yet to appear. In this light, we are not able to compare our use of explicit Stickelberger generators to their work. However, this talk does neither mention a short basis of the Stickelberger lattice, which is at the heart of our work, nor lift all obstructions to apply it to *any* conductor.

In December 2021, a "filtered-$\mathcal{S}$-unit software" was released by Bernstein, treating the prime $p \leq 43$ conductor case, on a webpage[5] describing the "simplest $\mathcal{S}$-unit attack" using a technique described in [BL21]. This work is not related to our construction. Finally, the authors of [BL21] argued that "spherical models" should not be applied to log-$\mathcal{S}$-unit lattices, which may have particular geometric properties. This phenomenon was experimentally observed already in [BR20], and is confirmed by all of our experiments in medium dimensions.

2 Preliminaries

Notations. For any $i, j \in \mathbb{Z}$ with $i \leq j$, the set of all integers between i and j is denoted by $[\![i, j]\!]$. For any $x \in \mathbb{Q}$, let $\{x\}$ denote its fractional part, i.e., such that $0 \leq \{x\} < 1$ and $x - \{x\} \in \mathbb{Z}$. A vector is represented by a bold letter $\mathbf{v}$, and for any $p \in \mathbb{N}^* \cup \{\infty\}$, its ℓ_p-norm is written $\|\mathbf{v}\|_p$. The n-dimensional vector with all 1's is denoted by $\mathbf{1}_n$. All matrices are given using *row* vectors.

2.1 Cyclotomic fields

We denote the cyclotomic field of conductor m, $m \not\equiv 2 \mod 4$, by $K_m = \mathbb{Q}[\zeta_m]$, where ζ_m is a primitive m-th root of unity. It has degree $n = \varphi(m)$, its maximal order is $\mathcal{O}_{K_m} = \mathbb{Z}[\zeta_m]$ ([Was97, Th. 2.6]), and its discriminant is given precisely by $\Delta_{K_m} = (-1)^{\varphi(m)/2} \frac{m^{\varphi(m)}}{\prod_{p|m} p^{\varphi(m)/(p-1)}}$ ([Was97, Pr. 2.7]), which is of order n^n.

In this paper, we consider *any* conductor $m > 1$ of the general prime factorization $m = p_1^{e_1} p_2^{e_2} \cdots p_t^{e_t}$, $m \not\equiv 2 \mod 4$, and let $q_i = p_i^{e_i}$ for all $i \in [\![1, t]\!]$. In particular, m has exactly t distinct prime divisors. Let G_m denote the Galois group of K_m, which can be made explicit by ([Was97, Th. 2.5]):

$$G_m = \left\{\sigma_s : \zeta_m \longmapsto \zeta_m^s;\ 0 < s < m, (s, m) = 1\right\} \simeq \left(\mathbb{Z}/m\mathbb{Z}\right)^\times.$$

[4] The slides are available at https://cr.yp.to/talks.html#2021.08.20.
[5] This is hosted by https://s-unit.attacks.cr.yp.to/filtered.html.

In particular, we denote by $\sigma_s \in G_m$ the automorphism sending any m-th root of unity to its s-th power. For convenience, the automorphism induced by complex conjugation is written $\tau = \sigma_{-1}$. The algebraic norm of $\alpha \in K_m$ is defined by $\mathcal{N}(\alpha) = \prod_{\sigma \in G_m} \sigma(\alpha)$, hence the absolute norm element in the integral group ring $\mathbb{Z}[G_m]$ is $N_m = \sum_{\sigma \in G_m} \sigma$.

Maximal real subfield. The maximal real subfield of K_m, written K_m^+, is the fixed subfield of K_m under complex conjugation, i.e., $K_m^+ := K_m^{\langle \tau \rangle} = \mathbb{Q}\big(\zeta_m + \zeta_m^{-1}\big)$. Its maximal order is $\mathcal{O}_{K_m^+} = \mathbb{Z}\big[\zeta_m + \zeta_m^{-1}\big]$ (see e.g. [Was97, Pr. 2.16]).

By Galois theory, since $\langle \tau \rangle$ is a normal subgroup of G_m, the maximal real subfield of K_m is a Galois extension of $\mathbb{Q}$ with Galois group $G_m^+ := \mathrm{Gal}\big(K_m^+/\mathbb{Q}\big)$ isomorphic to $G_m/\langle \tau \rangle$. We identify G_m^+ with the following system of representatives modulo τ restricted to K_m^+: $G_m^+ = \big\{\sigma_s|_{K_m^+};\ 0 < s < \frac{m}{2}, (s,m) = 1\big\}$. Technically, each $\sigma_s|_{K_m^+} \in G_m^+$ extends in G_m to either σ_s or $\tau\sigma_s = \sigma_{-s}$. For simplicity, we always choose to lift $\sigma_s|_{K_m^+} \in G_m^+$ to $\sigma_s \in G_m$ and drop the restriction to K_m^+ which should be clear from the context. This slight abuse of notation appears to be very practical. For example, the corestriction $\mathrm{Cor}_{K_m/K_m^+}\big(\sigma_s|_{K_m^+}\big)$, defined as the sum of all elements of G_m that restricts to $\sigma_s|_{K_m^+}$, namely $\sigma_s + \tau\sigma_s$, is written using the much simpler expression $(1+\tau)\cdot\sigma_s$.

2.2 Real and relative class groups

Fractional ideals of K_m form a multiplicative group $\mathcal{I}_m$ containing the normal subgroup $\mathcal{P}_m := \big\{\langle\alpha\rangle;\ \alpha \in K_m\big\}$ of principal ideals. The quotient group $\mathcal{I}_m/\mathcal{P}_m$ is called the *class group* of K_m and denoted by Cl_m. It is finite and its cardinal h_m is the *class number* of K_m. For any $\mathfrak{b} \in \mathcal{I}_m$, the class of $\mathfrak{b}$ in Cl_m is written $\big[\mathfrak{b}\big]$.

The integral group ring $\mathbb{Z}[G_m]$ acts naturally on $\mathcal{I}_m$; more precisely, for any element $\alpha = \sum_{\sigma \in G_m} a_\sigma \sigma \in \mathbb{Z}[G_m]$, and any $\mathfrak{b} \in \mathcal{I}_m$, $\mathfrak{b}^\alpha := \prod_{\sigma \in G_m} \sigma\big(\mathfrak{b}\big)^{a_\sigma}$. The class group and class number of the maximal real subfield K_m^+ are denoted respectively by Cl_m^+ and h_m^+. The relative norm map $\mathcal{N}_{K_m/K_m^+}$ induces a homomorphism from Cl_m to Cl_m^+, whose kernel is the so-called *relative class group*, written Cl_m^- and of cardinal the *relative class number* h_m^-. Hence, by construction, for any $\mathfrak{b}$ s.t. $\big[\mathfrak{b}\big] \in \mathrm{Cl}_m^-$, $\mathfrak{b}^{1+\tau} \cap K_m^+$ is principal. One important specificity of cyclotomic fields is that the real class group Cl_m^+ embeds into Cl_m *via* the natural inclusion map, which to each ideal class $\big[\mathfrak{b}\big] \in \mathrm{Cl}_m^+$ associates the ideal class $\big[\mathfrak{b}\cdot\mathcal{O}_{K_m}\big] \in \mathrm{Cl}_m$ [Was97, Th. 4.14]. Concretely, it implies that $h_m = h_m^+ \cdot h_m^-$ is the product of the plus part and the relative part of the class number.

Plus part and relative part of the class number. Generally, not much is known about the class number of a number field, and the analytic class number formula [Neu99, Cor. 5.11(ii)] allows obtaining a rough upper bound $h_m \leq \widetilde{O}\big(\sqrt{|\Delta_{K_m}|}\big)$.

In the case of cyclotomic fields though, the structure of the relative class group is better understood. Using analytic means, the relative class number has the following explicit expression [Was97, Th. 4.17]: $h_m^- = Qw \cdot \prod_{\chi \text{ odd}} \Big(-\frac{1}{2} B_{1,\chi}\Big)$,

where $w = 2m$ if m is odd and $w = m$ if m is even, $Q = 1$ if m is a prime power and $Q = 2$ otherwise, and $B_{1,\chi}$ is defined by $\frac{1}{f}\sum_{a=1}^{f} a\cdot\chi(a)$ for any odd primitive character χ modulo m of conductor f dividing m. Computing this value is in practice very efficient, using adequate representations of Dirichlet characters.

The hard part of cyclotomic class numbers computations is to obtain the plus part h_m^+, and relatively few of them are known. We use the values from [Was97, Tab. §4], [Mil14, Th. 1.1 and 1.2] and [BFHP21, Tab. 1], consistently assuming the *Generalized Riemann Hypothesis* (GRH). We also provide 58 additional values of h_m^+ in [BLNR21, Tab. 2.1] for completeness.

The fact that the plus part of the class number seems much smaller than the relative part is striking. Weber's conjecture claims that $h_{2^e}^+ = 1$ for any $e > 1$, and Buhler, Pomerance and Robertson [BPR04] argue, based on Cohen-Lenstra heuristics, that for all but finitely many pairs (p, e), where p is a prime and e is a positive integer, $h_{p^{e+1}}^+ = h_{p^e}^+$. For prime power conductors, this conjecture claims that the plus part is asymptotically constant. These conjectures are backed up by Schoof's extensive calculations [Sch03] in the prime conductor case, and by the above explicit values. In particular, under GRH, Miller proved Weber's conjecture up to $m = 512$, and we note that according to Schoof's table, $h_m^+ \leq \sqrt{m}$ holds for more than 96.6% of all prime conductors $m = p < 10000$.

Prime ideal classes generators. When picking a set of prime ideals in the algorithms of this paper, an important feature is that they generate the class group. In general, even assuming GRH, only a large bound on the norm of generators is known, indeed Bach proved [Bac90, Th. 4] that $\mathcal{N}(\mathfrak{L}_{\max}) \leq 12\ln^2|\Delta_{K_m}|$, where $\mathfrak{L}_{\max}$ is the biggest ideal inside a generating set of Cl_m of minimum norm. In practice though, this bound seems very pessimistic [BDF08, §6].

On the other hand, as prime ideals belong to Cl_m^- only with probability roughly $1/h_m^+$, searching for generators of the *subgroup* Cl_m^- mechanically increases the provable upper bound on generators. More precisely, writing as $\mathfrak{L}_{\max}^-$ the biggest ideal of a generating set of Cl_m^-, Wesolowski proved [Wes18, Rem. 2] that $\mathcal{N}(\mathfrak{L}_{\max}^-) \leq \left(2.71h_m^+ \cdot \ln|\Delta_{K_m}| + 4.13\right)^2$.

Finally, we use the notation $h_{m,(\mathfrak{L}_1,\ldots,\mathfrak{L}_k)}$ to denote the cardinal of the subgroup of Cl_m generated by the k classes $\left[\mathfrak{L}_i\right]$, i.e., the determinant of the kernel of $\mathfrak{f}_{\mathfrak{L}_1,\ldots,\mathfrak{L}_k} : (e_1,\ldots,e_k) \in \mathbb{Z}^k \longmapsto \prod_{1\leq i\leq k}\left[\mathfrak{L}_i\right]^{e_i} \in \mathrm{Cl}_m$.

2.3 Logarithmic $\mathcal{S}$-embeddings

We introduce log-$\mathcal{S}$-unit lattices and discuss proper normalization by the Product Formula that was at the heart of the practical improvements of [BR20] compared to [PHS19].

Places of the cyclotomic field K_m are usually split into two parts: the set $\mathcal{S}_\infty$ of *infinite* places can be identified with the (complex) embeddings of K_m into $\mathbb{C}$, up to conjugation; the set $\mathcal{S}_0$ of *finite* places is specified by the infinite set of prime ideals of K_m, each prime ideal $\mathfrak{p}$ inducing an embedding of K_m into its $\mathfrak{p}$-adic completion $K_{m,\mathfrak{p}}$. Hence, any place $v \in \mathcal{S}_\infty \cup \mathcal{S}_0$ induces an absolute value $|\cdot|_v$

on K_m, and Ostrowski's theorem for number fields [Nar04, Th. 3.3] shows that all possible absolute values on K_m are obtained in this way. Concretely, for $\alpha \in K_m$: $\forall \sigma \in \mathcal{S}_\infty, |\alpha|_\sigma = \left|\sigma(\alpha)\right|$ and $\forall \mathfrak{p} \in \mathcal{S}_0, |\alpha|_\mathfrak{p} = p^{-v_\mathfrak{p}(\alpha)}$, where $v_\mathfrak{p}(\cdot)$ is the valuation of α at $\mathfrak{p}$ and $\langle p \rangle = \mathfrak{p} \cap \mathbb{Z}$. A remarkable fact is that all these absolute values are tied by the *Product Formula* [Nar04, Th. 3.5]:

$$\forall \alpha \in K_m, \qquad \prod_{v \in \mathcal{S}_\infty \cup \mathcal{S}_0} |\alpha|_v^{[K_{m,v}:\mathbb{Q}_v]} = 1. \tag{2.1}$$

The $\mathcal{S}_\infty$-part of this product is $\left|\mathcal{N}(\alpha)\right|$, as for $\sigma \in \mathcal{S}_\infty$, $K_{m,\sigma} = \mathbb{C}$ and $\mathbb{Q}_\sigma = \mathbb{R}$, so that $[K_{m,\sigma} : \mathbb{Q}_\sigma] = 2$. Similarly, for $\mathfrak{p} \in \mathcal{S}_0$, we have $|\alpha|_\mathfrak{p}^{[K_{m,\mathfrak{p}}:\mathbb{Q}_p]} = \mathcal{N}(\mathfrak{p})^{-v_\mathfrak{p}(\alpha)}$.

$\mathcal{S}$-unit group structure. Fix a finite set $\mathcal{S}$ of places; in this paper we shall consider that $\mathcal{S}$ *always* contains $\mathcal{S}_\infty$. The so-called $\mathcal{S}$-unit group of K_m, denoted by $\mathcal{O}_{K_m,\mathcal{S}}^\times$, is the multiplicative subgroup of K_m generated by all elements whose valuations are non zero only at the finite places of $\mathcal{S}$. Formally:

$$\mathcal{O}_{K_m,\mathcal{S}}^\times = \left\{ \alpha \in K_m;\ \langle \alpha \rangle = \prod_{\mathfrak{p} \in \mathcal{S} \cap \mathcal{S}_0} \mathfrak{p}^{v_\mathfrak{p}(\alpha)} \right\}.$$

Theorem 2.1 (Dirichlet-Chevalley-Hasse [Nar04, Th. III.3.12, Cor. 1]). *The $\mathcal{S}$-unit group is the direct product of the group of roots of unity $\mu\big(\mathcal{O}_{K_m}^\times\big)$ and a free abelian group with $|\mathcal{S}| - 1$ generators. There exists a fundamental system of $\mathcal{S}$-units $\varepsilon_1, \ldots, \varepsilon_{|\mathcal{S}|-1}$ such that any $\varepsilon \in \mathcal{O}_{K_m,\mathcal{S}}^\times$ is uniquely written as: $\varepsilon = \mu \cdot \prod_{i=1}^{|\mathcal{S}|-1} \varepsilon_i^{k_i}$, where $\mu \in \langle \pm\zeta_m \rangle$ is a root of unity and $k_i \in \mathbb{Z}$.*

Log-$\mathcal{S}$-unit lattice. A fundamental ingredient of the proof of this theorem is to build an embedding of $\mathcal{O}_{K_m,\mathcal{S}}^\times$ into the real space of dimension $|\mathcal{S}|$, whose kernel is $\mu\big(\mathcal{O}_{K_m}^\times\big)$ and whose image is a lattice of dimension $\big(|\mathcal{S}|-1\big)$. This embedding is called the *logarithmic $\mathcal{S}$-embedding*, and its image is called the log-$\mathcal{S}$-unit lattice.

Several equivalent definitions of this logarithmic $\mathcal{S}$-embedding are acceptable for the proof. However, for cryptanalytic purposes, experimental evidence [BR20] suggests that it is crucial to use a properly normalized embedding for the decodability of the log-$\mathcal{S}$-unit lattice. Thus, we define [Nar04, §3, p.98]: $\mathrm{Log}_\mathcal{S}\, \alpha = \big([K_{m,v} : \mathbb{Q}_v] \cdot \ln|\alpha|_v\big)_{v \in \mathcal{S}} = \Big(\big\{2\ln|\sigma(\alpha)|\big\}_{\sigma \in \mathcal{S}_\infty}, \big\{-v_\mathfrak{p}(\alpha) \ln \mathcal{N}(\mathfrak{p})\big\}_{\mathfrak{p} \in \mathcal{S} \cap \mathcal{S}_0}\Big)$. By definition of $\mathcal{O}_{K_m,\mathcal{S}}^\times$, $\mathbb{R} \otimes \mathrm{Log}_\mathcal{S}\, \mathcal{O}_{K_m,\mathcal{S}}^\times$ is included in the hyperplane orthogonal to $\mathbf{1}_{|\mathcal{S}|}$. Showing that its dimension is at least $|\mathcal{S}| - 1$ is more involved.

A basis of the log-$\mathcal{S}$-unit lattice is given by the images $\mathrm{Log}_\mathcal{S}\, \varepsilon_i$ of the fundamental system of $\mathcal{S}$-units of Th. 2.1, as in [BR20, Eq. (2.7)]. Actually, we shall use later that for any maximal set of independent $\mathcal{S}$-units, their images under any logarithmic $\mathcal{S}$-embedding form a full rank sublattice of the corresponding log-$\mathcal{S}$-unit lattice. Its volume is given by [BR20, Pr. 2.2 and Eq. (2.8)].

As mentioned in [PHS19,BDPW20,BR20], a convenient trick in the context of the cryptanalysis of id-Svp is to consider an *expanded* version of the logarithmic $\mathcal{S}$-embedding, halving and repeating twice $\mathcal{S}_\infty$-coordinates, namely: $\overline{\mathrm{Log}}_\mathcal{S}\, \alpha = \Big(\big\{\ln|\sigma(\alpha)|, \ln|\sigma(\alpha)|\big\}_{\sigma \in \mathcal{S}_\infty}, \big\{[K_{m,\mathfrak{p}} : \mathbb{Q}_p] \cdot \ln|\alpha|_\mathfrak{p}\big\}_{\mathfrak{p} \in \mathcal{S} \cap \mathcal{S}_0}\Big)$.

In particular, this reduces the volume of the log-$\mathcal{S}$-unit lattice, as shown by [BR20, Pr. 2.3]. In practice though, we did not observe any fundamental difference between the approximation factors obtained using $\mathrm{Log}_{\mathcal{S}}$ or $\overline{\mathrm{Log}}_{\mathcal{S}}$.

2.4 Hard problems in Number Theory

One of the most difficult classical steps of the Approx-id-SVP algorithms proposed in [CDW17,PHS19,BR20,CDW21] is to find a solution to the ClDL defined as:

Problem 2.2 (Class Group Discrete Logarithm (ClDL)). Given a basis of prime ideals $\{\mathfrak{L}_1, \ldots, \mathfrak{L}_k\}$, and a challenge ideal $\mathfrak{b}$, find $\alpha \in K_m$ and integers $e_1, \ldots, e_k$ such that $\langle \alpha \rangle = \mathfrak{b} \cdot \prod_i \mathfrak{L}_i^{e_i}$, if this decomposition exists.

In this definition, we also ask for an *explicit* element α of the field, contrary to the definition of, e.g., [CDW17, Pr. 2]. Nevertheless, we note that in both quantum and classical worlds, the standard way to solve this problem boils down to computing $\mathcal{S}$-units, for $\mathcal{S}$ containing $\mathfrak{b}$ and the $\mathfrak{L}_i$'s, so that this explicit element is a byproduct of the resolution. Furthermore, put in this form it encompasses the well-known *Principal Ideal Problem* (PIP), using an empty set of ideals.

The *Shortest Generator Problem* (SGP) asks, from a generator α of a principal ideal, for the shortest generator α' such that $\langle \alpha \rangle = \langle \alpha' \rangle$. Similarly, we define:

Problem 2.3 (Shortest Class Group Discrete Logarithm (S-ClDL)). Given a solution $\langle \alpha \rangle = \mathfrak{b} \cdot \prod_i \mathfrak{L}_i^{e_i}$ to the ClDL problem, find $w_1, \ldots, w_k \in \mathbb{Z}_{\geq 0}$ and $\alpha' \in K_m$ such that $\langle \alpha' \rangle = \mathfrak{b} \cdot \prod_i \mathfrak{L}_i^{w_i}$ and α' is the smallest possible one.

The condition for the w_i's to be positive is crucial. Note that all recent algorithms for Approx-id-SVP that are not bound to principal ideals eventually output an approximate solution of the S-ClDL [CDW21,PHS19,BR20]. If the set of prime ideals is sufficiently large compared to $\mathfrak{b}$, then S-ClDL is exactly id-SVP.

We also mention the Close Principal Multiple (CPM) problem which, given an ideal $\mathfrak{b}$, asks to find $\mathfrak{c}$ such that $\mathfrak{bc}$ is principal and $\mathcal{N}(\mathfrak{c})$ is small. This specific problem is used in [CDW21], and the authors prove that under GRH and using a factor base containing all prime ideals of norm up to $m^{4+o(1)}$, there exists a solution $\mathfrak{c}$ with $\mathcal{N}(\mathfrak{c}) \leq \exp\big(\widetilde{O}(m^{1+o(1)})\big)$ [CDW21, §1.3.4].

Complexities. As shown in [BS16], class groups, unit groups, class group discrete logarithms and principal ideal generator computations can be reduced to $\mathcal{S}$-unit groups computations for appropriate sets of places $\mathcal{S}$. Denote by $\mathrm{T}_{\mathcal{S}}(K_m)$ the running time of the computation of the $\mathcal{S}$-unit group in K_m. Under GRH, in a quantum setting, $\mathrm{T}_{\mathcal{S}}(K_m) = \mathrm{poly}\big(\ln|\Delta_{K_m}|, \big|\mathcal{S}\big|, \max_{\mathfrak{p}\in\mathcal{S}} \ln \mathcal{N}(\mathfrak{p})\big)$ by [EHKS14,BS16]. In a classical setting, $\mathrm{T}_{\mathcal{S}}(K_m) = \mathrm{poly}\big(\big|\mathcal{S}\big|, \max_{\mathfrak{p}\in\mathcal{S}} \ln \mathcal{N}(\mathfrak{p})\big) \cdot \exp \widetilde{O}\big(\ln^{2/3}(|\Delta_K|\big)$ is mainly subexponential in the degree of the cyclotomic field K_m [BF14,PHS19]. The exponent can be lowered to $1/2$ when m is a prime power [BEF+17].

2.5 Lattices

Let L be a Euclidean lattice of full rank n. The first minimum $\lambda_1(L)$ of L is defined as the ℓ_2-norm of the smallest vector $\mathbf{v} \in L^*$, and the ℓ_2-distance from $\mathbf{t}$ to L, for any $\mathbf{t}$ in the span $L \otimes \mathbb{R}$ of L, is defined by $\mathrm{dist}_2(L, \mathbf{t}) = \min_{\mathbf{v} \in L} \|\mathbf{t} - \mathbf{v}\|_2$.

The *Approximate Shortest Vector Problem* (Approx-SVP) is, given a lattice L and an approximation factor af, to find $\mathbf{v} \in L$ such that $\|\mathbf{v}\|_2 \leq \mathrm{af} \cdot \lambda_1(L)$. Similarly, the *Approximate Closest Vector Problem* (Approx-CVP) asks, given a lattice L, an approximation factor af and a target $\mathbf{t}$ in the span $L \otimes \mathbb{R}$ of L, for a vector $\mathbf{v} \in L$ such that $\|\mathbf{t} - \mathbf{v}\|_2 \leq \mathrm{af} \cdot \mathrm{dist}_2(L, \mathbf{t})$. A practical Approx-CVP oracle is given by Babai's Nearest Plane algorithm [Bab86].

Bounding approximation factors. An ideal lattice of K_m is the full-rank image under the Minkowski embedding in $\mathbb{R}^{\varphi(m)}$ of a fractional ideal $\mathfrak{b}$ of K_m. Unlike generic lattices, a lower bound of the first minimum is implied by the arithmetic-geometric mean inequality, using that for any $b \in \mathfrak{b}$, $\mathcal{N}(\mathfrak{b})$ divides $|\mathcal{N}(b)|$. Thus:

$$\sqrt{n} \cdot \mathcal{N}(\mathfrak{b})^{1/n} \leq \lambda_1(\mathfrak{b}) \leq \sqrt{n} \cdot \mathcal{N}(\mathfrak{b})^{1/n} \sqrt{|\Delta_{K_m}|}^{1/n}, \tag{2.2}$$

where $n = \varphi(m) = \deg K_m$ and the right inequality is Minkowski's inequality. Actually, applying the Gaussian Heuristic to ideal lattices would give that on average, $\lambda_1(\mathfrak{b}) \approx \sqrt{\frac{n}{2\pi e}} \cdot \mathrm{Vol}^{1/n}(\mathfrak{b})$, where $\mathrm{Vol}(\mathfrak{b}) = \mathcal{N}(\mathfrak{b})\sqrt{|\Delta_{K_m}|}$. This hypothesis is commonly used for the analysis of cryptosystems based on structured lattices, and we note that the *exact* approximation factors reached by the Twisted-PHS algorithm in [BR20] match this heuristic.

For any $\mathbf{x} \in \mathfrak{b}$, let $\mathrm{af}(\mathbf{x}) = \|\mathbf{x}\|_2 / \lambda_1(\mathfrak{b})$ denote the approximation factor reached by $\mathbf{x}$ for the SVP in the ideal lattice $\mathfrak{b}$. In general, $\lambda_1(\mathfrak{b})$ is not known, but Eq. (2.2) imply the bounds $\mathrm{af}_{\mathsf{inf}}(\mathbf{x}) \leq \mathrm{af}(\mathbf{x}) \approx \mathrm{af}_{\mathsf{gh}}(\mathbf{x}) \leq \mathrm{af}_{\mathsf{sup}}(\mathbf{x})$, where:

$$\begin{gathered} \mathrm{af}_{\mathsf{inf}}(\mathbf{x}) := \frac{\|\mathbf{x}\|_2}{\sqrt{n} \cdot \mathrm{Vol}^{1/n}(\mathfrak{b})}, \qquad \mathrm{af}_{\mathsf{sup}}(\mathbf{x}) := \frac{\|\mathbf{x}\|_2}{\sqrt{n} \cdot \mathcal{N}(\mathfrak{b})^{1/n}}, \\ \mathrm{af}_{\mathsf{gh}}(\mathbf{x}) := \sqrt{2\pi e} \cdot \mathrm{af}_{\mathsf{inf}}(\mathbf{x}). \end{gathered} \tag{2.3}$$

Quality of a lattice basis. Several indicators have been used in the literature to attempt to measure the quality of a lattice basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ relatively to the SVP or the CVP. We will focus on the following three standard quantities:

1. the root-Hermite Factor $\delta_0(B)$, defined by $\delta_0^n(B) = \|\mathbf{b}_1\|_2 / \mathrm{Vol}^{1/n} B$, is commonly used to compare lattice reduction algorithms like LLL [LLL82] or BKZ [CN11]. On average, LLL reaches $\delta_0 \approx 1.022$ [GN08] whereas BKZ with blocksize $b \geq 50$ heuristically yields $\delta_0 \approx \left(\frac{b}{2\pi e}(\pi b)^{1/b}\right)^{1/(2b-2)}$ [Che13].
2. the (normalized) orthogonality defect $\delta(B)$, given by $\delta^n(B) = \prod_i \left(\frac{\|\mathbf{b}_i\|_2}{\mathrm{Vol}^{1/n} B}\right)$ [MG02, Def. 7.5] involves all vectors of the basis. By Minkowski's second theorem, its smallest possible value is upper bounded by $\sqrt{1 + \frac{n}{4}}$.
3. the logarithms of the norms of *Gram-Schmidt Orthogonalization* (GSO) vectors $\mathbf{b}_i^\star$ give also valuable information. For example, a rapid decrease in the sequence $\ln \|\mathbf{b}_i^\star\|_2$ at $i \geq 2$ indicates that $\mathbf{b}_i$ is rather not orthogonal to the previously generated subspace $\langle \mathbf{b}_1, \ldots, \mathbf{b}_{i-1} \rangle$.

3 An explicit full-rank family of independent $\mathcal{S}$-units

In this section, we exhibit a full rank family of *independent* $\mathcal{S}$-units, where the finite places $\mathcal{S}$ correspond to a collection of full Galois orbits of split prime ideals. As mentioned in introduction, this family is composed of three parts:

1. Circular units are recalled in §3.1 using the material from [Kuč92, Th. 6.1];
2. Stickelberger generators are in §3.2, sticking to the exposition of [BK21];
3. Real $\mathcal{S}^+$-units (apart from real units), where $\mathcal{S}^+$ is the set $\mathcal{S} \cap K_m^+$ of places of $\mathcal{S}$ restricted to K_m^+, are in §3.3.

Considering real $\mathcal{S}^+$-units and proving in §3.4 the multiplicative index of our family in the full $\mathcal{S}$-unit group constitute our main theoretical contributions. Finally, the saturation process used to mitigate this index is described in §3.5.

Remark 3.1. Recall that m has prime factorization $m = q_1 q_2 \cdots q_t \not\equiv 2 \mod 4$, where $q_i = p_i^{e_i} > 2$ for $i \in [\![1, t]\!]$. In this section, we will use subsets M_m^+ and M'_m of $[\![1, m]\!]$ that are useful to describe resp. a fundamental family of circular units and a short $\mathbb{Z}$-basis of the Stickelberger ideal of K_m. Their precise definitions from resp. [Kuč92, p.293] and [BK21, Eq. (11)] can be found in [BLNR21, §A.1].

3.1 Circular units

Circular units are sometimes called *cyclotomic units* in the literature, as in [Was97, §8]. We prefer to use the historical terminology from algebraic number theory, e.g. Sinnott [Sin78, §4] and Kučera [Kuč92, §2], in order to avoid any confusion with the whole unit group $\mathcal{O}_{K_m}^\times$ of the m-th cyclotomic field.

Definition 3.2 (Circular units [Was97, §8.1]). *Let V_m be the multiplicative subgroup of $K_m^\times$ generated by $\{1 - \zeta_m^a;\ 1 \le a \le m\}$. The group of* circular units *is the intersection $C_m := V_m \cap \mathcal{O}_{K_m}^\times$.*

Note that V_m contains the torsion of K_m, since $-\zeta_m = (1 - \zeta_m)/(1 - \zeta_m^{-1})$. The circular units form a subgroup of $\mathcal{O}_{K_m}^\times$ of finite index, more precisely:

Proposition 3.3 ([Sin78, Th. p.107]). *The index of C_m in $\mathcal{O}_{K_m}^\times$ is finite:*

$$\left[\mathcal{O}_{K_m}^\times : C_m\right] = 2^b \cdot h_m^+, \qquad \text{with } b = \begin{cases} 0 & \text{if } t = 1, \\ 2^{t-2} + 1 - t & \text{otherwise,} \end{cases}$$

where t is the number of distinct prime divisors of m.

Hence, circular units provide a very large subgroup of $\mathcal{O}_{K_m}^\times$: indeed, the real part of the class number is expected to be small (§2.2), and the other factor *generically* grows linearly in m (see [HW38, Th. 430 and 431] for a precise statement).

An explicit system of fundamental circular units for any m has been given in [GK89] and independently in [Kuč92, Th. 6.1]. More precisely, for $0 < a < m$, define the following special circular units, where $m_i = m/p_i^{e_i}$ [Kuč92, p.176]:

$$v_a = \begin{cases} 1 - \zeta_m^a & \text{if } \forall i \in [\![1, t]\!],\ m_i \nmid a, \\ \dfrac{1 - \zeta_m^a}{1 - \zeta_m^{m_i}} & \text{otherwise, for the unique } m_i \mid a. \end{cases} \tag{3.1}$$

Theorem 3.4 ([Kuč92, Th. 6.1]). *Recall the definition of $M_m^+ \subsetneq [\![1, m]\!]$ can be found in [BLNR21, §A.1]. The set $\{v_a;\ a \in M_m^+\}$ is a system of fundamental circular units of K_m: for any circular unit $\eta \in C_m$, there exist a uniquely determined map $k : M_m^+ \to \mathbb{Z}$, and a root of unity $\mu \in \langle \pm\zeta_m \rangle$ s.t. $\eta = \mu \cdot \prod_{a \in M_m^+} v_a^{k(a)}$.*

A crucial point for the cryptanalysis of id-SVP in [CDW21] is that the logarithmic embedding of these elements is short. Namely, explicitly writing the constants that appear in the proof of [CDW21, Lem. 3.5], we have, for any $0 < a < m$, that $\|\mathrm{Log}_{\mathcal{S}_\infty}(1 - \zeta_m^a)\|_2 \leq 1.32 \cdot \sqrt{m}$.

3.2 Stickelberger generators

In this section, we use [BK21, Th. 3.1] to describe a short *basis* of the so-called Stickelberger ideal, viewed as a $\mathbb{Z}$-module. These Stickelberger short relations correspond to principal ideals whose generators are surprisingly easy to compute using Jacobi sums as in [BK21, §6]. Following Sinnott [Sin80], for all $a \in \mathbb{Z}$, let:

$$\theta_m(a) = \sum_{s \in (\mathbb{Z}/m\mathbb{Z})^\times} \left\{-\frac{as}{m}\right\} \cdot \sigma_s^{-1} \in \mathbb{Q}\big[G_m\big], \tag{3.2}$$

and let N_m be the absolute norm element $N_m = \sum_{\sigma \in G_m} \sigma$.

Definition 3.5 (Stickelberger ideal [Sin80, p.189]). *Let $\mathcal{S}'_m$ be the $\mathbb{Z}$-module of $\mathbb{Q}\big[G_m\big]$ generated by $\{\theta_m(a);\ 0 < a < m\} \cup \{\frac{1}{2}N_m\}$. The* Stickelberger ideal *of K_m is the intersection $\mathcal{S}_m = \mathcal{S}'_m \cap \mathbb{Z}\big[G_m\big]$.*

As in [CDW21], we shall refer to the *Stickelberger lattice* when $\mathcal{S}_m$ is viewed as a $\mathbb{Z}$-module. Note that in some references, like in [Was97, §6.2], the Stickelberger ideal is defined as the smaller ideal $\mathbb{Z}\big[G_m\big] \cap \theta_m(-1)\mathbb{Z}\big[G_m\big]$, which coincides with Def. 3.5 if and only if m is a prime power [Kuč86, Pr. 4.3].

Theorem 3.6 (Stickelberger's theorem [Sin80, Th. 3.1]). *The Stickelberger ideal $\mathcal{S}_m$ of K_m annihilates the class group of K_m. Hence, for any ideal $\mathfrak{b}$ of K_m and any $\alpha = \sum_{\sigma \in G_m} a_\sigma \sigma \in \mathcal{S}_m$, the ideal $\mathfrak{b}^\alpha = \prod_{\sigma \in G_m} \sigma(\mathfrak{b})^{a_\sigma}$ is principal.*

An outstanding point is that the proof of this important result is completely explicit, i.e., for any $\alpha \in \mathcal{S}_m$, and any fractional ideal $\mathfrak{b}$ of K_m, an explicit $\gamma \in K_m$ s.t. $\langle \gamma \rangle = \mathfrak{b}^\alpha$ is constructed. It appears that when α is a short element of $\mathcal{S}_m$, this explicit generator is very efficiently computable.

A short basis of the Stickelberger lattice. An element of the integral group ring $\mathbb{Z}\big[G_m\big]$ is called *short* if it is of the form $\sum_{\sigma \in G_m} a_\sigma \sigma \in \mathbb{Z}\big[G_m\big]$, where $a_\sigma \in \{0, 1\}$ for all $\sigma \in G_m$. Short elements of $\mathcal{S}_m$ have been identified in [Sch08, Th. 9.3(i) and Ex. 9.3] in the prime conductor case, and the proof has been adapted to any conductor in [CDW21, Lem. 4.4] to prove the shortness of the following generating set of $\mathcal{S}_m$:

$$W = \{w_a;\ a \in [\![2, m]\!]\}, \quad \text{with } w_a = \theta_m(1) + \theta_m(a-1) - \theta_m(a). \tag{3.3}$$

Note that using $\theta_m(a) + \theta_m(-a) = N_m$ when $m \nmid a$, we obtain $w_a = w_{m-a+1}$ whenever $1 < a < m$, and that $w_m = N_m$ using also $\theta_m(m) = 0$. Hence, W is the set $\{w_a;\ 2 \leq a \leq \lceil \frac{m}{2} \rceil\} \cup \{N_m\}$.

We emphasize that only knowing a generating set of short elements as in [CDW21] is not necessarily sufficient. Though it would be possible to build a basis from this generating set to solve the CVP like in [CDW21, Cor. 2.2], without any geometric loss using e.g. [MG02, Lem. 7.1], we observed that the slight euclidean norm growth of the obtained basis vectors translates into a dramatic increase of the size of the (possibly rational) coefficients of the corresponding generators, in a way that significantly hinders subsequent computations. In particular, in order to climb dimensions as far as possible and best approach log-$\mathcal{S}$-unit lattices using the saturation process described in §3.5, it is crucial to constrain both the number of elements we use and their size, i.e., to use a *basis* of the Stickelberger lattice containing only *short* elements. In [BK21], a very large family of short elements [BK21, Pr. 3.1] encompassing $W \setminus \{N_m\}$ is made explicit:

Proposition 3.7 ([BK21, Pr. 3.1]). *Let $a, b \in \mathbb{Z}$ satisfying $m \nmid a$, $m \nmid b$ and $m \nmid (a+b)$. Then $\alpha = \theta_m(a) + \theta_m(b) - \theta_m(a+b)$ is a short element of $\mathcal{S}_m$. Moreover, $(1+\tau)\cdot\alpha = N_m$, so exactly one half of the coefficients of α are zeros.*

Then, from this family, a short basis is computationally easy to extract:

Theorem 3.8 ([BK21, Th. 3.6]). *Recall $M'_m \subsetneq [\![1, m]\!]$ is defined in [BLNR21, §A.1]. There exists an efficiently computable map $\alpha_m(\cdot)$ from $[\![1, m]\!]$ to the family of short elements of $\mathcal{S}_m$ described in Pr. 3.7, s.t. $\{\alpha_m(c);\ c \in M'_m\} \cup \{N_m\}$ is a $\mathbb{Z}$-basis of the Stickelberger lattice $\mathcal{S}_m$ of K_m having only short elements.*

The explicit definition of $\alpha_m(\cdot)$ is given in [BK21, §3.2], and included for completeness in [BLNR21, §A.2]. We stress that when m is a prime, this basis coincides with the one given by [Sch08, Th. 9.3(i)] and with the set W in Eq. (3.3).

Effective Stickelberger generators using Jacobi sums. As previously mentioned, the proof of Th. 3.6 is explicit, i.e., for any $\alpha \in \mathcal{S}_m$ and any fractional ideal $\mathfrak{b}$ of K_m, it builds an explicit $\gamma \in K_m$ s.t. $\langle\gamma\rangle = \mathfrak{b}^\alpha$ [Was97, §6.2], [Sin80, §3.1]. Moreover, when α is a short basis element from Th. 3.8, it turns out that γ has a simple expression using Jacobi sums [BK21, §5].

We briefly treat the split case here. Let $\ell \in \mathbb{Z}$ be a prime s.t. $\ell \equiv 1 \mod m$, and let $\mathfrak{L}$ be any fixed (split) prime ideal of K_m above ℓ. Let a, b be such as in Pr. 3.7, then for $\alpha = \theta_m(a) + \theta_m(b) - \theta_m(a+b)$, we have that $\mathfrak{L}^\alpha$ is a principal ideal generated by the following Jacobi sum [BK21, Pr. 5.1]:

$$\mathcal{J}_\mathfrak{L}(a,b) = -\sum_{u \in \mathcal{O}_{K_m}/\mathfrak{L}} \chi_\mathfrak{L}^a(u)\chi_\mathfrak{L}^b(1-u) \quad \in K_m, \tag{3.4}$$

where $\chi_\mathfrak{L}(u) \in \langle\zeta_m\rangle$ verifies $\chi_\mathfrak{L}(u) \equiv u^{(\ell-1)/m} \mod \mathfrak{L}$, for any $u \in (\mathcal{O}_{K_m}/\mathfrak{L})^\times$, and $\chi_\mathfrak{L}(0) = 0$. When $\alpha = \alpha_m(c)$ for $c \in M'_m$, we shall write $\gamma_{\mathfrak{L},c}$ for the generator of $\mathfrak{L}^{\alpha_m(c)}$. Using a discrete logarithm table for elements of $(\mathcal{O}_{K_m}/\mathfrak{L})^\times$, the computation, for a fixed prime $\mathfrak{L}$, of all Jacobi sums corresponding to the short basis $\{\alpha_m(c);\ c \in M'_m\}$ is very fast.

3.3 Real $\mathcal{S}^+$-units

A consequence of Th. 3.8, since $\left|M'_m\right| = \frac{\varphi(m)}{2}$, is that the Stickelberger lattice only has rank $\frac{\varphi(m)}{2}+1$ in $\mathbb{Z}\left[G_m\right]$; in particular, it is not full rank, hence cannot be directly used as a lattice of class relations. In previous works, obtaining a full rank lattice in $\mathbb{Z}[G_m]$ from $\mathcal{S}_m$ was done by projecting into $(1-\tau)\mathcal{S}_m$ [CDW21, §4.3], or by the adjunction of $(1+\tau)\mathbb{Z}[G_m]$ [CDW17, Def. 2]. Both can be used as a lattice of class relations for the *relative* class group Cl_m^-. In particular, the so-called *augmented* Stickelberger lattice $\mathcal{S}_m+(1+\tau)\mathbb{Z}[G_m]$ annihilates the relative class group and has full rank in $\mathbb{Z}[G_m]$, as shown in [CDW17, Lem. 2].

We generalize this result by considering the module of all real class group relations between relative norm ideals of ideals from the entire class group Cl_m. In §3.4, we shall prove that the Stickelberger lattice augmented with these real class group relations yields a lattice of class relations for the *whole* class group. Note that, as opposed to other modules like $(1-\tau)\mathcal{S}_m$ or $\mathcal{S}_m+(1+\tau)\mathbb{Z}\left[G_m\right]$, real class group relations actually depend on the underlying prime ideals.

On one hand, this affects negatively the shortness of the obtained relation vectors: putting those in Hermite Normal Form, we shall see later that each relation, viewed as a vector of integer valuations, has ℓ_2-norm at most h_m^+. On the other hand, removing the constraint to belong to the relative class group brings a significant practical and theoretical gap: first, it allows choosing prime ideals of smallest possible norms, which as shown in [BR20, §3.3] or [CDW21, Th. 4.8] lowers in practice the obtained approximation factor; second, whereas prime ideals of norm at most Bach's bound are sufficient to generate the entire class group, prime generators for the *relative* class group are only proven to be of norm bounded by the *larger* bound $(2.71\cdot h_m^+\cdot\ln\Delta_{K_m}+4.13)^2$ from [Wes18].

Lifting real class group relations. Let $\ell_1,\ldots,\ell_d$ be distinct prime integers satisfying $\ell_i\equiv 1 \mod m$, so that ℓ_i splits in K_m, for all i in $[\![1,d]\!]$. For each i, fix a prime ideal $\mathfrak{L}_i \mid \ell_i$ in K_m of norm ℓ_i, and let $\mathfrak{l}_i=\mathcal{N}_{K_m/K_m^+}\left(\mathfrak{L}_i\right)=\mathfrak{L}_i^{1+\tau}\cap K_m^+$ be the relative norm ideal of $\mathfrak{L}_i$. Since $\mathfrak{L}_i$ is a split prime ideal of K_m dividing ℓ_i, the ideal $\mathfrak{l}_i$ is a split prime ideal of K_m^+ of norm ℓ_i, and by Kummer-Dedekind's theorem we have $\mathfrak{l}_i\cdot\mathcal{O}_{K_m}=\mathfrak{L}_i^{1+\tau}$. This justifies the slight abuse of notation of writing $\mathfrak{l}_i^\sigma=\mathfrak{L}_i^{(1+\tau)\sigma}\cap K_m^+$, for any $\sigma\in G_m$.

We are interested in the real class group relations between all prime ideals in the G_m^+-orbits of the $\mathfrak{l}_i$, i.e., between the following prime ideals of K_m^+:

$$\left\{\mathfrak{l}_i^{\sigma_s};\ i\in[\![1,d]\!], 0<s<\tfrac{m}{2}, (s,m)=1\right\}. \tag{3.5}$$

The important point is, any class relation in K_m^+ between ideals from Eq. (3.5) translates to a class relation in K_m using repeatedly $\mathfrak{l}_i^\sigma\cdot\mathcal{O}_{K_m}=\mathfrak{L}_i^{(1+\tau)\sigma}$. More precisely, let $\left(r_1,\ldots,r_d\right)\in\mathbb{Z}\left[G_m^+\right]^d$ represent a real class relation in K_m^+ between ideals $\left\{\mathfrak{l}_i^{\sigma_s}\right\}$ of Eq. (3.5), i.e., there exists $\gamma_r^+\in K_m^+$ s.t. $\gamma_r^+\cdot\mathcal{O}_{K_m^+}=\prod_{i=1}^d\mathfrak{l}_i^{r_i}$. Then, this relation lifts naturally to a class relation $\left((1+\tau)\cdot r_1,\ldots,(1+\tau)\cdot r_d\right)$ in K_m between prime ideals in the G_m-orbits $\left\{\mathfrak{L}_i^\sigma;\ i\in[\![1,d]\!],\sigma\in G_m\right\}$ as:

$$\gamma_r^+\cdot\mathcal{O}_{K_m}=\prod_{i=1}^d\mathfrak{L}_i^{(1+\tau)r_i}. \tag{3.6}$$

Let $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ denote the lattice of class relations between elements of all G_m^+-orbits of $\{\mathfrak{l}_i;\ i \in [\![1,d]\!]\}$. Concretely, it is the kernel of the following map:

$$\mathfrak{f}_{\mathfrak{l}_1,\dots,\mathfrak{l}_d} : (r_{i,s})_{\substack{1\le i\le d,\\ 0<s<m/2,(s,m)=1}} \in \mathbb{Z}^{d\cdot\frac{\varphi(m)}{2}} \longmapsto \prod_{i,s} \left[\mathfrak{l}_i^{\sigma_s}\right]^{r_{i,s}} \in \mathrm{Cl}_m^+ . \tag{3.7}$$

Using the canonical isomorphism of $\mathbb{Z}$-modules $\mathbb{Z}^{d\cdot\frac{\varphi(m)}{2}} \simeq_{\mathbb{Z}} \mathbb{Z}[G_m^+]^d$, the lattice of class relations $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ may be viewed as a $\mathbb{Z}$-submodule of $\mathbb{Z}[G_m^+]^d$. Lifting all these relations back to K_m as in Eq. (3.6), we therefore obtain the submodule $(1+\tau)\cdot C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d} \subseteq (1+\tau)\mathbb{Z}[G_m]^d$, that we shall call the lattice of *real class relations* between the G_m-orbits of $\{\mathfrak{L}_i;\ i \in [\![1,d]\!]\}$.

Remark 3.9. When $h_m^+ = 1$, $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ is isomorphic to d copies of the integral group ring $\mathbb{Z}[G_m^+]$ and the lattice of real class relations is simply $(1+\tau)\mathbb{Z}[G_m]^d$.

Euclidean norm of real class relations. We now identify a real class group relation from $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ to a vector in $\mathbb{Z}^{d\cdot\frac{\varphi(m)}{2}}$. In other words, we consider only the valuations of these relations on the G_m^+-orbits of the prime ideals $\mathfrak{l}_1,\dots,\mathfrak{l}_d$. Furthermore, $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ is put in Hermite Normal Form, conveniently for the proof of the following proposition, provided in the full version of this paper [BLNR21], but better bounds might easily be obtained using e.g. the LLL algorithm.

Proposition 3.10. *Suppose the lattice $C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ of real class relations is in HNF. Then, for all $\mathbf{w} \in C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d} \subseteq \mathbb{Z}[G_m^+]^d$, we have $\|\mathbf{w}\|_2 \le \|\mathbf{w}\|_1 \le h_m^+$.*

This means that $(1+\tau)\cdot C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ can be used in the CDW algorithm instead of $(1+\tau)\mathbb{Z}[G_m]^d$, as we will see in §4, while still reaching the same asymptotic approximation factor under the same assumption on the Galois-module structure of Cl_m [CDW21, Ass. 1], as long as $h_m^+ \le O(\sqrt{m})$. This slightly more restrictive hypothesis (see the discussion in §2.2) will be more than compensated by the fact that it removes the need for the $\mathfrak{l}_i$'s to be principal, which has a significant impact in practice on the algebraic norm of the chosen ideals, and thus on the final approximation factor reached in [CDW21, Alg. 6].

Explicit real generators. For each relation $r = (r_1,\dots,r_d) \in C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$, we compute an explicit $\gamma_r^+ \in K_m^+ \subsetneq K_m$ that verifies Eq. (3.6). Together with the unit group $\mathcal{O}_{K_m^+}^\times$ of K_m^+, they form a fundamental system of $\mathcal{S}^+$-units, where the finite places of $\mathcal{S}^+$ are the G_m^+-orbits of the relative norm ideals $\mathfrak{l}_i$.

In the next section, we shall see that adding the explicit Stickelberger generators of §3.2 to these real generators yields a maximal set of independent $\mathcal{S}$-units in the degree $\varphi(m)$ cyclotomic field K_m, at the much smaller cost of computing a fundamental system of real $\mathcal{S}^+$-units in K_m^+ of degree only $\frac{\varphi(m)}{2}$.

In practice, though this remains the main bottleneck of our experimental setting, it allows us to push effectively our experiments up to degree $\varphi(m) = 210$, whereas the full $\mathcal{S}$-unit group computations of [BR20] were bound to $\varphi(m) = 70$.

3.4 A $\mathcal{S}$-unit subgroup of finite index

As in §3.3, let $\ell_1, \dots, \ell_d$ be prime integers satisfying $\ell_i \equiv 1 \mod m$; for each i, fix a (split) prime ideal $\mathfrak{L}_i \mid \ell_i$ in K_m and let $\mathfrak{l}_i = \mathfrak{L}_i \cap K_m^+$. Let $\mathcal{S}$ be a set of places containing, apart from the infinite places of K_m, all G_m-orbits of the $\mathfrak{L}_i$'s. Combining the results of §3.1, §3.2 and §3.3, we get the following family of $\mathcal{S}$-units:

$$\mathfrak{F} = \left\{v_a;\ a \in M_m^+\right\} \cup \left\{\gamma^-_{\mathfrak{L}_i,b};\ i \in [\![1,d]\!], b \in M'_m\right\} \cup \left\{\gamma^+_r;\ r \in C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}\right\} \quad (3.8)$$

where the first set is the set of *circular units* given by Th. 3.4, the second is the set of explicit *Stickelberger generators* stated at the end of §3.2 and the last one is the set of *real generators* as in Eq. (3.6).

This family has $\big(\varphi(m)/2-1\big)+d\cdot\varphi(m)$ elements, which matches precisely the multiplicative rank of the full $\mathcal{S}$-unit group modulo torsion $\mathcal{O}^\times_{K_m,\mathcal{S}}/\mu\big(\mathcal{O}^\times_{K_m}\big)$.[6] In this section, we prove that these $\mathcal{S}$-units are indeed independent and we compute the index of the subgroup of $\mathcal{O}^\times_{K_m,\mathcal{S}}$ generated by those elements.

Theorem 3.11. *Let $h_{m,(\mathfrak{L}_1,\dots,\mathfrak{L}_d)}$ (resp. $h^+_{m,(\mathfrak{l}_1,\dots,\mathfrak{l}_d)}$) be the cardinal of the subgroup of Cl_m (resp. Cl^+_m) generated by the G_m-orbits of $\mathfrak{L}_1,\dots,\mathfrak{L}_d$ (resp. the G^+_m-orbits of $\mathfrak{l}_1,\dots,\mathfrak{l}_d$). The family $\mathfrak{F}$ given in Eq. (3.8) is a maximal set of independent $\mathcal{S}$-units. The subgroup generated by $\mathfrak{F}$ in $\mathcal{O}^\times_{K_m,\mathcal{S}}/\mu\big(\mathcal{O}^\times_{K_m}\big)$ has index:*

$$\left(\frac{h_m \cdot h^+_{m,(\mathfrak{l}_1,\dots,\mathfrak{l}_d)}}{h_{m,(\mathfrak{L}_1,\dots,\mathfrak{L}_d)}}\right) \cdot 2^b \cdot \big(h^-_m\big)^{d-1} \cdot \left(2^{\frac{\varphi(m)}{2}-1} \cdot 2^a\right)^d,$$

where $a = b = 0$ if m is a prime power, and $a = 2^{t-2}-1$, $b = 2^{t-2}+1-t$ when m has t distinct prime divisors.

When the G_m-orbits of the $\mathfrak{L}_i$'s generate Cl_m, the first term in this index equals h^+_m. As we shall see in §3.5, the powers of 2 can be killed by saturation techniques, so the problem comes from the $(h^-_m)^{d-1}$ part, which has generically *huge* prime factors. Intuitively, this is because the Stickelberger relations miss all class group relations that exist between two (or more) distinct G_m-orbits.

First, we show that the lattice obtained by adding one copy of the Stickelberger ideal per G_m-orbit, to the lattice $(1+\tau)\cdot C^+_{\mathfrak{l}_1,\dots,\mathfrak{l}_d}$ of real class relations, yields a full-rank submodule of $\mathbb{Z}[G_m]^d$. Hence, we have obtained a full-rank lattice of class relations for the union of all G_m-orbits above $\ell_1,\dots,\ell_d$.

We begin by restricting our attention to the case $d = 1$. We need the following lemma, which extends and proves an observation already made in [DPW19, Rem. 3] in the prime conductor case (see [BLNR21, §3.4] for the full proofs):

Lemma 3.12. *The index of $\mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G^+_m]$ in $\mathbb{Z}[G_m]$ is finite:*

$$\big[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot\mathbb{Z}[G^+_m]\big] = 2^{\varphi(m)/2-1} \cdot 2^a \cdot h^-_m,$$

where $a = 0$ if $t = 1$ and $a = 2^{t-2}-1$ else, where m has t prime divisors.

[6] Note that for our purpose, the torsion units play no role and can thus be put aside.

When $h_m^+ = 1$, the lattice of real class relations is always $(1+\tau)\cdot\mathbb{Z}[G_m^+]$, and Lem. 3.12 gives the whole story. In the general case $h_m^+ \neq 1$, we deduce:

Lemma 3.13. *Let ℓ be a prime integer that splits in K_m, let $\mathfrak{L} \mid \ell$ in K_m and let $\mathfrak{l} = \mathfrak{L}^{1+\tau} \cap K_m^+$. Let $h_{m,(\mathfrak{l})}^+$ be the cardinal of the subgroup of Cl_m^+ generated by the G_m^+-orbit of $\mathfrak{l}$ in K_m^+. The $\mathbb{Z}$-module generated by $\mathcal{S}_m$ and the lattice $(1+\tau)\cdot C_{\mathfrak{l}}^+$ of real class relations of the G_m-orbit of $\mathfrak{L}$, has finite index in $\mathbb{Z}[G_m]$:*

$$\left[\mathbb{Z}[G_m] : \mathcal{S}_m + (1+\tau)\cdot C_{\mathfrak{l}}^+\right] = 2^{\varphi(m)/2-1}\cdot 2^a \cdot h_m^- \cdot h_{m,(\mathfrak{l})}^+,$$

where $a = 0$ if $t = 1$ and $a = 2^{t-2}-1$ else, where m has t prime divisors.

Finally, for the case where there are $d \geq 1$ orbits, a reasoning very similar to the proofs of Lem. 3.12 and 3.13 leads to:

Proposition 3.14. *Let $h_{m,(\mathfrak{l}_1,\ldots,\mathfrak{l}_d)}^+$ be the cardinal of the subgroup of Cl_m^+ generated by all G_m^+-orbits of $\mathfrak{l}_1,\ldots,\mathfrak{l}_d$. Then, the $\mathbb{Z}$-module generated by the lattice $(1+\tau)\cdot C_{\mathfrak{l}_1,\ldots,\mathfrak{l}_d}^+ \subseteq (1+\tau)\cdot\mathbb{Z}[G_m^+]^d$ of real class relations between the G_m-orbits of the $\mathfrak{L}_i$'s, and the diagonal block matrix of d copies of $\left(\mathcal{S}_m \setminus N_m\mathbb{Z}\right)$, verifies:*

$$\left[\mathbb{Z}[G_m]^d : \mathcal{S}_m^d + (1+\tau)\cdot C_{\mathfrak{l}_1,\ldots,\mathfrak{l}_d}^+\right] = \left(2^{\varphi(m)/2-1}\cdot 2^a \cdot h_m^-\right)^d \cdot h_{m,(\mathfrak{l}_1,\ldots,\mathfrak{l}_d)}^+.$$

Proof of Th. 3.11. The independence comes from Pr. 3.14 and the trivial fact that circular units are independent from Stickelberger and real generators. The index of the subgroup generated by $\mathfrak{F}$ in $\mathcal{O}_{K_m,\mathcal{S}}^\times/\mu\left(\mathcal{O}_{K_m}^\times\right)$ is given by:

$$\left[\mathcal{O}_{K_m}^\times : C_m\right] \cdot \frac{\left[\mathbb{Z}[G_m]^d : \mathcal{S}_m^d + (1+\tau)\cdot C_{\mathfrak{l}_1,\ldots,\mathfrak{l}_d}^+\right]}{\left|\det\left(\ker \mathfrak{f}_{\mathcal{S}}\right)\right|},$$

where $\ker \mathfrak{f}_{\mathcal{S}}$ is the lattice of all class group relations between finite places of $\mathcal{S}$. The first term is given by Pr. 3.3 and the numerator of the second term by Pr. 3.14. By definition of $\mathcal{O}_{K_m,\mathcal{S}}^\times$, the denominator is precisely $h_{m,(\mathfrak{L}_1,\ldots,\mathfrak{L}_d)}$. Rearranging terms adequately yields the result. □

3.5 Saturation

Saturation is a standard tool of computational algebraic number theory that has been used in various contexts like unit and class group computations, and can be traced back at least to [PZ89, §5.7]. This procedure is described in more detail in [BLNR21, §3.5], and we refer to e.g. [BFHP21, §4.3] for a formal exposition.

Intuitively, the e-saturation procedure applied to $\mathfrak{F}$ consists in detecting e-th powers in the subgroup generated by $\mathfrak{F}$, including their e-th roots in the set and rebuilding a basis of multiplicatively independent elements. At the end, the index of the new basis is no longer divisible by e. Remark that the output size does not depend on e, but only on the number and size of the elements of $\mathfrak{F}$.

As the index given by Th. 3.11 is divisible by a large power of 2, it is therefore natural to 2-saturate $\mathfrak{F}$ in order to mitigate its exponential growth, obtaining the 2-saturated family $\mathfrak{F}_{\mathsf{sat}}$. Note however that the relative class number h_m^- in the index of Th. 3.11 hides *huge* prime factors that at first glance render this strategy hopeless in general to obtain the full $\mathcal{S}$-unit group from $\mathfrak{F}$.

4 Removing quantum steps from the CDW algorithm

The full material for this section is given in [BLNR21, §B], we summarize the main points here. The CDW algorithm for solving Approx-SVP was introduced in [CDW17] for cyclotomic fields of prime power conductors, using short relations of the Stickelberger lattice as a keystone. [CDW21] extended it to all conductors.

In this section, we show how to use the results of §3.2, §3.3 and §3.4 to remove most quantum steps of [CDW21]. More precisely, we first propose an equivalent rewriting of [CDW21, Alg. 7] that enlightens some hidden steps that reveal useful for subsequent modifications. Then, we plug in the explicit generators of §3.2 ([BK21]) and Eq. (3.6) for relative class group orbits, to remove the last call to the quantum PIP solver. Finally, by considering the module of *all* real class group relations, using Pr. 3.14 and Th. 3.11, we remove the need of a random walk mapping any ideal of K_m into Cl_m^-, at the (small) additional price of restricting to cyclotomic fields such that $h_m^+ \leq O(\sqrt{m})$.

An equivalent rewriting of CDW [BLNR21, §B.2]. Omitting details, the CDW algorithm works as follows, for any challenge ideal $\mathfrak{a}$ of K_m [CDW21, Alg. 7]:

1. Random walk to Cl_m^-: find $\mathfrak{b}$ such that $[\mathfrak{ab}] \in \mathrm{Cl}_m^-$.
2. Solve the ClDL of $\mathfrak{ab}$ on G_m-orbits of the prime ideals $\mathfrak{L}_1, \ldots, \mathfrak{L}_d$ of Cl_m^-. This gives a vector[7] $\epsilon = (\epsilon_1, \ldots, \epsilon_d) \in \mathbb{Z}[G_m]^d$ such that $\mathfrak{ab} \cdot \prod_i \mathfrak{L}_i^{\epsilon_i}$ is principal.
3. Solve the CPMP by projecting each ϵ_i in $\pi(\mathcal{S}_m) = (1-\tau)\mathcal{S}_m$, find a close vector $v_i = y_i \cdot \pi(\mathcal{S}_m)$ and lift v_i to get some η_i s.t. $\pi(\eta_i) = v_i$, $\|\epsilon - \eta\|_1$ is small *with positive coordinates*, and $\mathfrak{ab} \cdot \prod_i \mathfrak{L}_i^{\epsilon_i - \eta_i}$ is principal.
4. Apply the PIP algorithm of [BS16] to get a generator of this principal ideal.
5. Reduce the obtained generator by circular units like in [CDPR16].

This eventually outputs $h \in \mathfrak{a}$ of length $\|h\|_2 \leq \exp\big(\widetilde{O}(\sqrt{m})\big) \cdot \mathcal{N}(\mathfrak{a})^{1/\varphi(m)}$.

We focus on the lift procedure of Step 3. In [CDW21], $v \in \pi(\mathcal{S}_m)$ is lifted to $\eta \in \mathcal{S}_m$ with non-negative coordinates by setting $(\eta_\sigma, \eta_{\tau\sigma}) = (v_\sigma, 0)$ if $v_\sigma \geq 0$ and $(0, -v_\sigma)$ otherwise, for all $\sigma \in G_m^+$. This works because $[\mathfrak{c}]^{-1} = [\mathfrak{c}^\tau]$ for any $\mathfrak{c} \in \mathrm{Cl}_m^-$, but hides which exact product of relative norm ideals is involved. We propose a totally equivalent lift procedure: from $v = y \cdot \pi(\mathcal{S}_m)$, consider the preimage $\tilde{\eta} = y \cdot \mathcal{S}_m$. Define η by removing $\min\{\tilde{\eta}_\sigma, \tilde{\eta}_{\tau\sigma}\}$ to each $\tilde{\eta}_\sigma$ coordinate. Now, it is obvious that η is a combination y of relations in $\mathcal{S}_m$, and of relative norm relations given by the min part. Details are in [BLNR21, Alg. B.6].

Using explicit Stickelberger generators [BLNR21, §B.3]. Each element w_a of the generating set W of $\mathcal{S}_m$ corresponds to a generator $\mathcal{J}_{\mathfrak{L}}(1, a-1)$ (see §3.2). Similarly, each relative norm ideal writes $\langle \gamma_s^+ \rangle = \mathfrak{L}^{(1+\tau)\sigma_s}$ (see §3.3). Hence, from an (explicit) ClDL solution $\langle \alpha \rangle = \mathfrak{ab} \cdot \mathfrak{L}^{\epsilon}$, and given a CPMP solution, explicitly written as above as $\eta = y \cdot W + u \cdot (1+\tau) \cdot \mathbb{Z}[G_m^+]$, we have that a generator of $\mathfrak{ab} \cdot \mathfrak{L}^{\epsilon - \eta}$ is directly given by $\alpha / \big(\prod_a \mathcal{J}_{\mathfrak{L}}(1, a-1)^{y_a} \prod_s (\gamma_s^+)^{u_s} \big)$. This allows us to remove the quantum PIP in dimension n in step 4 (for each query). In exchange, we need to compute (only once) all real generators for relative norm relations, which can be done in dimension $\varphi(m)/2$ by [BS16, Alg. 2].

[7] In the CDW algorithm, the explicit generator given by the ClDL solver is discarded.

Avoiding the random walk [BLNR21, §B.4]. Finally, note that several quantum steps are performed (for each query) in the random walk that maps ideals to Cl_m^-. Using the results of §3.3, we replace the module $(1+\tau)\cdot\mathbb{Z}[G_m]^d$ by the module of all real class group relations. Asymptotically, we prove in [BLNR21, Pr. B.7] that this does not change the bound on the approximation factor obtained in [CDW21, Th. 5.1], under the same assumption on the Galois-module structure of Cl_m [CDW21, Ass. 1], as long as we restrict to fields K_m with $h_m^+ \leq O(\sqrt{m})$. This additional tiny assumption is largely compensated by the fact that only two quantum steps remain: one is performed only once in dimension $\varphi(m)/2$ to compute real class group relations and generators, and the second is solving the ClDL for each query (see [BLNR21, Tab. B.1]).

5 Computing log-$\mathcal{S}$-unit sublattices in higher dimension

Our main goal is to simulate the Twisted-PHS algorithm for high degree cyclotomic fields. To this end, we compute full-rank sublattices of the full log-$\mathcal{S}$-unit lattice using the knowledge of the maximal set $\mathfrak{F}$ of independent $\mathcal{S}$-units defined by Eq. (3.8) and its 2-saturated counterpart $\mathfrak{F}_{\mathsf{sat}}$ from §3.5. These sets are lifted from a complete set of real $\mathcal{S}^+$-units (see §3.3), hence are obtained at the classically subexponential cost of working in the half degree maximal real subfield. We note that by Th. 3.11, the index of these families grows rapidly as the number of orbits increases, hence these approximated modes give an upper bound on the approximation factors that can be expected when using Twisted-PHS.

The Twisted-PHS algorithm is briefly recalled in §5.1, and our experimental setting is detailed in §5.2. Then, we analyse in §5.3 the geometric characteristics of our log-$\mathcal{S}$-unit sublattices and the obtained approximation factors in §5.4.

5.1 The Twisted-PHS algorithm

The Twisted-PHS algorithm [BR20] was introduced as an improvement of the PHS algorithm [PHS19]. Both aim at solving Approx-id-SVP in any number field and have the same theoretically proven bounds for running time and reached approximation factors. However, the explicit $\mathcal{S}$-units formalism in [BR20] leads to a proper normalization of the used log-$\mathcal{S}$-embedding, weighting coordinates according to finite places norms. This turned out to give experimentally significant improvements on the lattices' decodability and on reached approximation factors.

Both algorithms are split in a *preprocessing phase*, performed only once for a fixed number field, and a *query phase*, for each challenge ideal. More precisely:

1. The preprocessing phase consists in choosing a set of finite places $\mathcal{S}$ generating the class group, computing the corresponding log-$\mathcal{S}$-unit lattice for an appropriate log-$\mathcal{S}$-embedding, and preparing the lattice for subsequent Approx-CVP requests using the Laarhoven's algorithm from [Laa16];
2. For each challenge ideal $\mathfrak{b}$, the query phase consists in first solving the ClDL relatively to $\mathcal{S}$, obtaining $\langle\alpha\rangle = \mathfrak{b}\cdot\prod_{\mathfrak{L}\in\mathcal{S}}\mathfrak{L}^{v_\mathfrak{L}}$. Then, this element is projected onto the span of the above log-$\mathcal{S}$-unit lattice, and a close vector of this lattice

gives a $\mathcal{S}$-unit s s.t. α/s is hopefully small. Here, guaranteeing that $\alpha/s \in \mathfrak{b}$ is achieved by applying a drift parameterized by some β on the target.

In the Twisted-PHS case, since the obtained lattice, after proper normalization, appears to have exceptionally good geometric characteristics, it was proposed to replace Laarhoven's algorithm by a lazy BKZ reduction in the preprocessing phase and Babai's Nearest Plane algorithm in the query phase [BR20, Alg. 4.2 and 4.3]. We will consider only this practical version in our experiments.

In details, for a number field K, the log-$\mathcal{S}$-unit lattice used in the Twisted-PHS algorithm is defined as $\varphi_{\mathsf{tw}}(\mathcal{O}_{K,\mathcal{S}}^{\times})$, where φ_{tw} is the log-$\mathcal{S}$-embedding given by $f_H \circ \overline{\mathrm{Log}}_{\mathcal{S}}$ [BR20, Eq. (4.1)], for an isometry f_H from the span H of $\overline{\mathrm{Log}}_{\mathcal{S}}$ to $\mathbb{R}^k$, where k equals the multiplicative rank of $\mathcal{O}_{K,\mathcal{S}}^{\times}$ modulo torsion.

Among the consequences of the proper normalization induced by $\overline{\mathrm{Log}}_{\mathcal{S}}$, the authors showed how to optimally choose a set of finite places that generate the class group [BR20, Alg. 4.1]. Namely, taking ideals of increasing prime norms in the set $\mathcal{S}$, they noticed that the density of the associated (twisted) log-$\mathcal{S}$-unit lattice $\varphi_{\mathsf{tw}}(\mathcal{O}_{K,\mathcal{S}}^{\times})$ increases up to an optimal value before decreasing.

Finally, a tricky aspect of the resolution resides in guaranteeing that the output solution is indeed an element of the challenge ideal, i.e., that $v_{\mathfrak{L}}(\alpha/s) \geq 0$ for all $\mathfrak{L} \in \mathcal{S} \cap \mathcal{S}_0$. In [BR20], this is done by applying a drift vector in the span of the log-$\mathcal{S}$-unit lattice, parameterized by some β whose optimal value is searched using a dichotomic strategy in the query phase. Concretely [BR20, Eq. (4.7)]:

$$\mathbf{t} = f_H\left(\left\{\ln|\alpha|_\sigma - \frac{k\beta + \ln \mathcal{N}(\mathfrak{b}) - \sum_{\mathfrak{L}\in\mathcal{S}} \ln \mathcal{N}(\mathfrak{L})}{[K:\mathbb{Q}]}\right\}_\sigma, \left\{\ln|\alpha|_{\mathfrak{L}}^{[K_{\mathfrak{L}}:\mathbb{Q}_\ell]} + \beta - \ln \mathcal{N}(\mathfrak{L})\right\}_{\mathfrak{L}\in\mathcal{S}}\right).$$

5.2 Experimental settings

Computing the full group of $\mathcal{S}$-units in a classical way is rapidly intractable, even in the case of cyclotomic fields; therefore, experiments performed in [BR20] on Twisted-PHS were bound to $\varphi(m) \leq 70$. We apply the Twisted-PHS algorithm using our full-rank sublattices of the whole log-$\mathcal{S}$-unit lattice induced by the independent family $\mathfrak{F}$ of Eq. (3.8), its 2-saturated counterpart $\mathfrak{F}_{\mathsf{sat}}$ (§3.5) and, when feasible, a fundamental system $\mathfrak{F}_{\mathsf{su}}$ for the full $\mathcal{S}$-unit group. Approximated modes with $\mathfrak{F}$ or $\mathfrak{F}_{\mathsf{sat}}$ give a glimpse on how Twisted-PHS scales in higher dimensions, where asymptotic phenomena like the growth of h_m start to express.

Source code and hardware description. All experiments have been implemented using SageMath v9.0 [Sag20], except for the full $\mathcal{S}$-unit groups computations for which we used Magma [BCP97], which appears much faster for this particular task and also offers an indispensable product ("Raw") representation. Moreover, fplll [FpL16] was used to perform all lattice reduction algorithms. The entire source code is provided on `https://github.com/ob3rnard/Tw-Sti`.

Most of the computations were performed in less than two weeks on a server with 72 Intel® Xeon® E5-2695v4 @2.1GHz with 768GB of RAM, using 2TB of storage for the precomputations. Real class group computations were performed on a single Intel® Core™ i7-8650U @3.2GHz CPU using 10GB of RAM.

m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+	m	$\varphi(m)$	h_m^+
136	64	2	408	128	2	205	160	2	356	176	†	520	192	4	265	208	†
212	104	5	268	132	†	328	160	†	376	184	†	840	192	†	424	208	†
145	112	2	284	140	†	440	160	5	191	190	11	303	200	†	636	208	†
183	120	4	292	144	†	163	162	4	221	192	†	404	200	†			
248	120	4	504	144	4	332	164	†	388	192	†	309	204	†			
272	128	2	316	156	†	344	168	†	476	192	†	412	204	†			

Table 5.1 – List of ignored conductors (†: failure to compute Cl_m^+ within a day).

Targeted cyclotomic fields. We consider cyclotomic fields of *any* conductor m s.t. $20 < \varphi(m) \leq 210$ with known real class number $h_m^+ = 1$, including those from [BLNR21, Tab. 2.1]. The restriction to $h_m^+ = 1$ is only due to technical interface obstructions, i.e., we are not aware of how to access the non-trivial real class group relations internally computed by SageMath. Additionally, for some of the conductors, we were not able to obtain the real class group within a day. Thus, we are left with 210 distinct cyclotomics fields, and Tab. 5.1 lists all ignored conductors.

Finite places choice. The optimal set of places computed by [BR20, Alg. 4.1] yields a number $d_{\max}$ of split G_m-orbits of smallest norms maximizing the density of the corresponding full log-$\mathcal{S}$-unit lattice. However, the index of our log-$\mathcal{S}$-unit *sublattices*, given by Th. 3.11, grows too quickly, roughly in $(h_m^-)^{d-1}$, so that their density always decreases as soon as $d > 1$. This remark motivates us to compute all log-$\mathcal{S}$-unit sublattices for $d = 1$ to $d_{\max}$ first split G_m-orbits.

Full rank log-$\mathcal{S}$-unit sublattices. The first maximal set of independent $\mathcal{S}$-units that we consider is $\mathfrak{F}$ from Eq. (3.8). The 2-saturation process of §3.5 mitigates the huge index of $\mathfrak{F}$, yielding family $\mathfrak{F}_{\mathsf{sat}}$. A fundamental system $\mathfrak{F}_{\mathsf{su}}$ of the full $\mathcal{S}$-unit group $\mathcal{O}_{K_m,\mathcal{S}}^{\times}$ (modulo torsion) is also used whenever it is computable in reasonable time, i.e., up to $\varphi(m) < 80$. As noted in §2.3, their images under any log-$\mathcal{S}$-embedding φ form full-rank sublattices resp. L_{urs}, L_{sat}, L_{su}, generated by resp. $\varphi(\mathfrak{F})$, $\varphi(\mathfrak{F}_{\mathsf{sat}})$, $\varphi(\mathfrak{F}_{\mathsf{su}})$, of the corresponding full log-$\mathcal{S}$-unit lattice $\varphi(\mathcal{O}_{K_m,\mathcal{S}}^{\times})$.

We consider several choices of the log-$\mathcal{S}$-embedding φ. Namely, we tried to evaluate the advantage of using the expanded $\overline{\mathrm{Log}}_{\mathcal{S}}$ (exp) over $\mathrm{Log}_{\mathcal{S}}$, labelled tw (as twisted by $[\mathbb{C} : \mathbb{R}] = 2$). We also considered versions with (iso) or without (noiso) the isometry f_H of [BR20, Eq. (4.2)]. This yields four choices for φ, e.g. tag noiso/tw is $\varphi = \mathrm{Log}_{\mathcal{S}}$ and iso/exp gives the original $\varphi_{\mathsf{tw}} = f_H \circ \overline{\mathrm{Log}}_{\mathcal{S}}$.

Compact product representation. In order to avoid the exponential growth of algebraic integers viewed in $\mathbb{Z}[x]/\langle \Phi_m(x) \rangle$, we use a compact product representation, so that any element α in $\mathfrak{F}$ (resp. $\mathfrak{F}_{\mathsf{sat}}$ or $\mathfrak{F}_{\mathsf{su}}$) is written on a set $g_1, \ldots, g_N$ of N small elements as $\alpha = \prod_{i=1}^{N} g_i^{e_i}$. Hence, besides the g_i's, each α is stored as a vector $e \in \mathbb{Z}^N$, and for any choice of φ, we have $\varphi(\alpha) = \sum_{i=1}^{N} e_i \cdot \varphi(g_i)$. This allows us to compute φ without the coefficient explosion encountered in [BR20, §5], which unlocks the full log-$\mathcal{S}$-unit lattices computations beyond degree 60.

m	d	set	k	$\mathrm{Vol}^{1/k}$	δ			$\max_{1\le i\le k}\|\mathbf{b}_i\|_2$		
					raw	LLL	bkz_{40}	raw	LLL	bkz_{40}
152	1	urs	107	8.691	2.016	1.570	1.551	45.007	38.466	38.202
		sat	107	6.928	4.398	1.787	1.822	752.306	23.280	21.720
		su	107	6.928	28.396	1.805	1.828	3163.723	21.953	21.446
	2	urs	179	9.683	2.157	1.623	1.590	48.754	41.313	41.404
		sat	179	7.384	7.670	1.885	1.896	6273.562	23.280	22.772
		su	179	6.816	65.355	2.226	2.322	3427.134	23.221	24.741
211	1	urs	314	14.325	2.672	2.291	2.257	96.068	97.930	96.569
		sat	314	11.386	9.998	2.581	2.562	9742.552	59.387	59.578
	5	urs	1154	18.232	3.118	2.542	2.497	118.124	119.160	115.888
		sat	1154	13.341	19.443	2.918	2.901	32067.612	71.428	72.752
	7	urs	1574	18.976	3.161	2.557	2.512	120.838	121.129	119.020
		sat	1574	13.771	26.841	2.927	2.910	530646.708	71.428	72.752

Table 5.2 – Geometric characteristics of L_{urs}, L_{sat} and L_{su} for $\mathbb{Q}(\zeta_{152})$ and $\mathbb{Q}(\zeta_{211})$ with log-$\mathcal{S}$-embedding φ_{tw} (of type $\mathsf{iso}/\mathsf{exp}$). For *all* bases, the root-Hermite factor verifies $|\delta_0 - 1| < 10^{-3}$.

Lattice reductions. For each of the constructed log-$\mathcal{S}$-unit sublattices, i.e. for each number of orbits $d \in [\![1, d_{\max}]\!]$, for each family of independent $\mathcal{S}$-units $\mathfrak{F}$, $\mathfrak{F}_{\mathsf{sat}}$ and (when feasible) $\mathfrak{F}_{\mathsf{su}}$, and for each choice of log-$\mathcal{S}$-embedding, we compare several levels of reduction: no reduction ("raw"), LLL-reduction and BKZ_{40}-reduction.

5.3 Geometry of the lattices

For all described choices of log-$\mathcal{S}$-unit sublattices, we first evaluate several geometrical parameters (see §2.5): reduced volume $V^{1/k}$, root-Hermite factor δ_0, orthogonality defect δ. We only give here a few examples giving a glimpse of what happens in general, and additional data can be found in [BLNR21, §C.1].

Table 5.2 contains data for cyclotomic fields $\mathbb{Q}(\zeta_{152})$ and $\mathbb{Q}(\zeta_{211})$ of degrees resp. 72 and 210. All values correspond to the $\mathsf{iso}/\mathsf{exp}$ log-$\mathcal{S}$-embedding, i.e., $\varphi = \varphi_{\mathsf{tw}}$. Indeed, as illustrated by [BLNR21, Tab. C.2], we experimentally note that using $(\mathsf{no})\mathsf{iso}/\mathsf{exp}$ seems geometrically slightly better than using $(\mathsf{no})\mathsf{iso}/\mathsf{tw}$. Notice how small is the normalized orthogonality defect after only LLL reduction, unambiguously below the tight Minkowski bound $\sqrt{1+\frac{k}{4}}$.

We then look at the logarithm of the Gram-Schmidt norms, for every described choice of log-$\mathcal{S}$-unit sublattices. Figure 5.1 plots the Gram-Schmidt log norms before and after BKZ reduction of the lattices L_{sat}, using the original $\mathsf{iso}/\mathsf{exp}$ log-$\mathcal{S}$-embedding φ_{tw}. As in [BR20, Fig. B.1–10], for each field the two curves are almost superposed, which is consistent with the observations on the orthogonality defect. We also checked the impact of the log-$\mathcal{S}$-embedding choice among all four options on the Gram-Schmidt logarithm norms of the *unreduced* basis $\varphi(\mathfrak{F}_{\mathsf{sat}})$. As expected, the isometry f_H has no influence on the Gram-Schmidt norms. On the other hand, using $\mathrm{Log}_{\mathcal{S}}$ or $\overline{\mathrm{Log}}_{\mathcal{S}}$ seems to alter only the

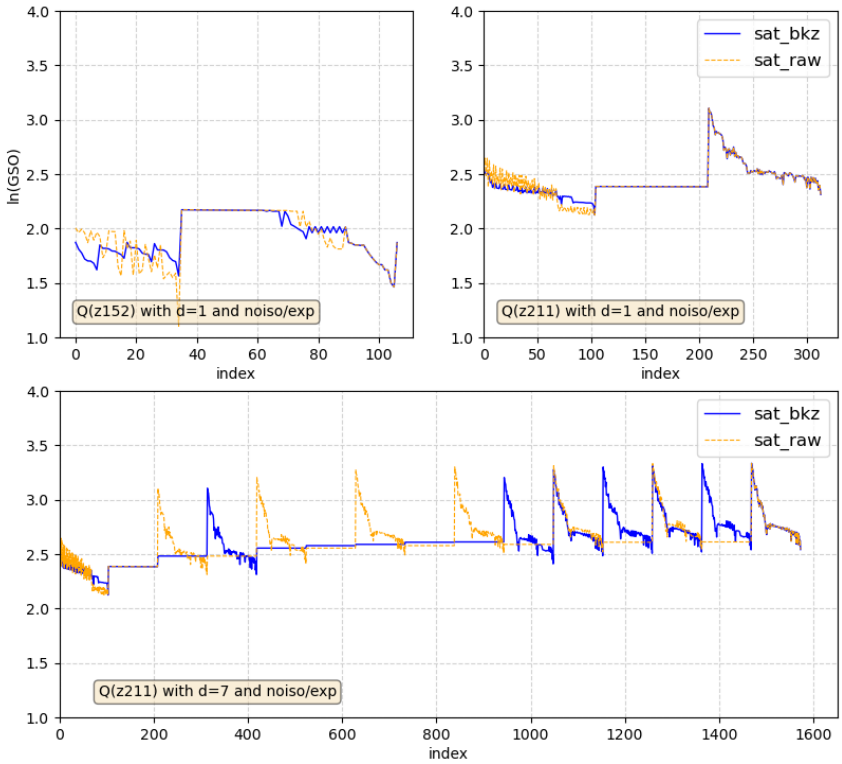

Fig. 5.1 – L_{sat} lattices for $\mathbb{Q}(\zeta_{152})$ and $\mathbb{Q}(\zeta_{211})$: Gram-Schmidt log norms before and after reduction by BKZ_{40}.

first norms, and in a very small way. This can be seen in [BLNR21, Tab. C.4]. Again, increasing the number of orbits does not influence these behaviours.

We stress that these very peculiar geometric characteristics – shape of the logarithm of the norms of the Gram-Schmidt basis, ease of reduction, very small orthogonality defect (after LLL) – already observed in [BR20, §5.1–2], are consistently viewed across all conductors, degrees, log-$\mathcal{S}$-unit sublattices and number of orbits. To give a concrete idea of e.g. the striking ease of reduction of these log-$\mathcal{S}$-unit sublattices, we report that for $m = 211$, BKZ_{40} terminates in around 7 minutes (resp. 30 minutes) on the log-$\mathcal{S}$-unit sublattice of dimension $k = 1154$ (resp. 1574) corresponding to $d = 5$ (resp. $d_{\max} = 7$), which is unusually fast.

This very broad phenomenon suggests that the explanation is possibly deep, an observation that has been recently developed by Bernstein and Lange [BL21].

5.4 Evaluation of the approximation factor

In [BR20], evaluating in practice the approximation factors reached by Twisted-PHS is done by choosing random split ideals of prime norm, solving the ClDL for these challenges and comparing the length of the obtained algebraic integer

with the length of the exact shortest element. As the degrees of the fields grow, solving the ClDL and exact id-Svp becomes rapidly intractable. Hence, we resort to simulating random outputs of the ClDL, similarly to [DPW19, Hyp. 8], and estimate the obtained approximation factors with inequalities from Eq. (2.3).

Simulation of ClDL *solutions.* To simulate targets that heuristically correspond to explicit generators α output by the ClDL, we assume that for each ideal $\mathfrak{L}_\mathsf{i} \in \mathcal{S}$, the vector $\left(v_{\mathfrak{L}_\mathsf{i}^\sigma}(\alpha)\right)_{\sigma\in G_m}$ of $\mathbb{Z}^{\frac{\varphi(m)}{2}}$ is uniform modulo the lattice of class relations, and that after projection along the $\mathbf{1}$-axis, $\left(\ln|\sigma(\alpha)|\right)_\sigma$ is uniform modulo the log-unit lattice. These hypotheses have already been used in [DPW19, Hyp. 8] or [BR20, H. 4.8], and are backed up by theoretical results in [BDPW20, Th. 3.3].

Drawing random elements modulo a lattice of rank k is done by following a Gaussian distribution of sufficiently large deviation. Concretely, we first choose a random split prime p in $[\![2^{97}, 2^{103}]\!]$. Then, for each $\mathfrak{L} \in \mathcal{S}\cap\mathcal{S}_0$, we pick random valuations $v_\mathfrak{L}(\alpha)$ modulo the lattice of class relations of rank $\left|\mathcal{S}\cap\mathcal{S}_0\right|$ and random elements $(u_\sigma)_{\sigma\in G_m^+} \in \mathbb{R}^{\varphi(m)/2}$ in the span of the log-unit lattice of rank $\frac{\varphi(m)}{2}-1$. Finally, we simulate $(\ln|\sigma(\alpha)|)_\sigma$ by adding $\frac{\ln p+\sum_{\mathfrak{L}\in\mathcal{S}} v_\mathfrak{L} \ln \mathcal{N}(\mathfrak{L})}{\varphi(m)}$ to each coordinate u_σ, so that their sum is $\frac{\ln|\mathcal{N}(\alpha)|}{2}$. For each field, we thereby generate 100 random targets on which to test Twisted-PHS on all lattice versions.

Reconstruction of a solution. For each simulated ClDL generator α, given as a random vector $(\{\ln|\sigma(\alpha)|\}_{\sigma\in G_m^+}, \{v_\mathfrak{L}(\alpha)\}_{\mathfrak{L}\in\mathcal{S}\cap\mathcal{S}_0})$, it is easy to compute $\varphi(\alpha)$ for any log-$\mathcal{S}$-embedding φ and to derive a target as in [BR20, Eq. (4.7)], including a drift parameterized by some β. Then, considering e.g. $L_\mathsf{sat} = \varphi(\mathfrak{F}_\mathsf{sat})$, given by the *BKZ*$_{40}$*-reduced* basis $U_\mathsf{bkz}\cdot\varphi(\mathfrak{F}_\mathsf{sat})$, we find a close vector $v = (y\cdot U_\mathsf{bkz})\cdot\varphi(\mathfrak{F}_\mathsf{sat})$ to this target using Babai's Nearest Plane algorithm, and from y, U_bkz and $\mathfrak{F}_\mathsf{sat}$ we easily recover, in compact representation, $s \in \mathcal{O}_{K_m,\mathcal{S}}^\times$ s.t. $v = \varphi(s)$ and also α/s.

The purpose of the drift parameter β is to guarantee $v_\mathfrak{L}(\alpha/s) \geq 0$ on all finite places. As mentioned in [BR20], the length of α/s is extremely sensitive to the value of β, so that they searched for an optimal value by dichotomy. However, this positiveness property actually does not seem to be monotonic in β, and in practice, using the same β on each finite place coordinate is too coarse when the dimension grows, resulting in unnecessarily large approximation factors. We instead obtained best results using random drifts in ℓ_∞-norm balls of radius 1 centered on the $\mathbf{1}$ axis. A first sampling of $O(\varphi(m))$ random points $\beta\cdot\mathbf{1}+\mathcal{B}_\infty(1)$ for a wide range of random β's allows us to select a β_0 around which we found the best $\|\alpha/s\|_2$ with all $v_\mathfrak{L}(\alpha/s)$ being positive. Then we sample $O(\varphi(m))$ uniform random points in the neighbourhood of β_0, namely in $[0.9\beta_0, 1.1\beta_0]\cdot\mathbf{1}+\mathcal{B}_\infty(1)$, and output the overall optimal $\|\alpha/s\|_2$ having all $v_\mathfrak{L}(\alpha/s) \geq 0$.

Estimator of the approximation factor. Since we do not have access to the shortest element of a challenge ideal, we cannot compute an exact approximation factor as in [BR20]. Instead, we estimate the retrieved approximation factor using the inequalities implied by Eq. (2.3). We focus on the Gaussian Heuristic, which gives consistent results with the exact approximation factors found in [BR20],

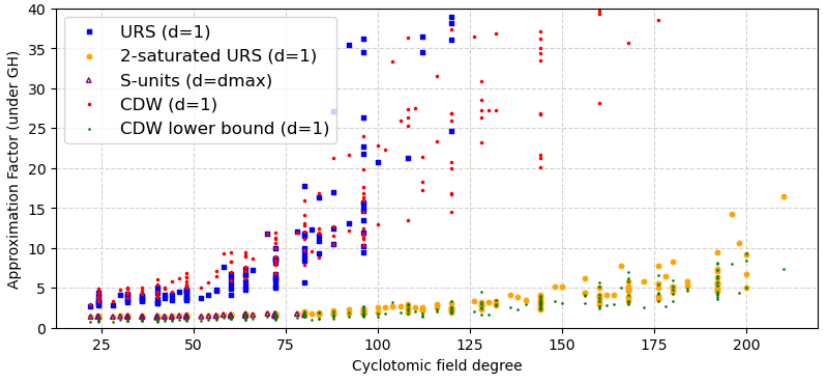

Fig. 5.2 – Approximation factors, with Gaussian Heuristic, reached by Tw-PHS for cyclotomic fields of degree up to 210, on lattices L_{urs}, L_{sat} and L_{su}.

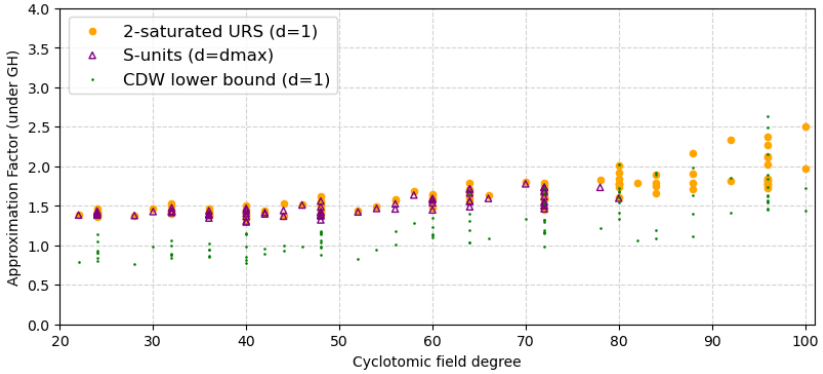

Fig. 5.3 – Approximation factors, with Gaussian Heuristic, reached by Tw-PHS for cyclotomic fields of degree up to 100, on lattices L_{sat} and L_{su}.

in small dimensions. For each cyclotomic field, the plotted points are the means, over the 100 simulated random targets, of the minimal approximation factors obtained using options iso/noiso and exp/tw. For each family $\mathfrak{F}$, $\mathfrak{F}_{\mathsf{sat}}$ and $\mathfrak{F}_{\mathsf{su}}$, we chose to keep only the factor base that gives the best result. This systematically translated into using $d = 1$ G_m-orbit for $\mathfrak{F}$ and $\mathfrak{F}_{\mathsf{sat}}$, whereas we had to use $d = d_{\max}$ for $\mathfrak{F}_{\mathsf{su}}$, as predicted by the Twisted-PHS algorithm.

Figure 5.2 shows the approximation factor $\mathrm{af}_{\mathsf{gh}}$ obtained for all lattices L_{urs}, L_{sat} and L_{su} (when feasible) after BKZ_{40} reduction. Figure 5.3 is a zoom of Fig. 5.2 that focuses on L_{sat} and L_{su} on small dimensions. First, we remark that using $\mathfrak{F}$ from Eq. (3.8), the retrieved approximation factors are increasing rapidly. Using the 2-saturated family $\mathfrak{F}_{\mathsf{sat}}$ yields much better results, and looking closely at Fig. 5.3 shows that using a basis $\mathfrak{F}_{\mathsf{su}}$ of the full $\mathcal{S}$-unit group, when feasible, even improves the picture if $d_{\max} > 1$, in which case L_{su} is denser than L_{sat}. For L_{su}, we stress that we obtain estimated approximation factors very similar to the exact ones observed in [BR20].

More generally, we observe a very strong correlation between the density of our lattices and the obtained approximation factors – the denser, the better. As an important related remark, the variance seen for $\mathrm{af}_{\mathsf{gh}}$ in Fig. 5.2 for distinct fields of same degree follows the variations of the norm of the first split prime, thus of the reduced volume of the considered log-$\mathcal{S}$-unit sublattice. We expect this variance to be smoothed through conductors for the full log-$\mathcal{S}$-unit lattice.

Furthermore, considering $m = 211$, the $\mathfrak{F}$ family gives $\mathrm{Vol}^{1/314} L_{\mathsf{urs}} \approx 14.325$ and an estimated $\mathrm{af}_{\mathsf{gh}} \approx 13170$, for $\mathfrak{F}_{\mathsf{sat}}$ we get $\mathrm{Vol}^{1/314} L_{\mathsf{sat}} \approx 11.386$ and a much smaller estimated $\mathrm{af}_{\mathsf{gh}} \approx 16.4$, whereas the optimal number of orbits predicted by the Twisted-PHS factor base choice algorithm [BR20, Alg. 4.1] is $d_{\max} = 7$, which yields a full log-$\mathcal{S}$-unit lattice of reduced volume only $\mathrm{Vol}^{1/1574} L_{\mathsf{su}} \approx 9.635$.

Comparison to the CDW algorithm. Using the same experimental setting, we compute the approximation factors obtained using the CDW algorithm as implemented in [DPW19] ("Naive version") with additional BKZ_{40} lattice reductions, as well as the experimentally derived *volumetric lower bound* from [DPW19, Eq. (5) and Tab. 1]. Those values are also represented in Fig. 5.2 and 5.3.

We note that our experimental results using the $\mathfrak{F}_{\mathsf{sat}}$ family are comparable to this volumetric lower bound. Moreover, for some fields, e.g. in dimensions 96, 160, 168, 200, this lower bound is defeated by the (approximated version of the) Twisted-PHS algorithm. Note that this does not invalidate the lower bound itself, which is stated for the two-phase CDW algorithm, but indicates the power of combining both steps in only one lattice as in the Twisted-PHS algorithm.

Acknowledgements. The first author is deeply indebted to Radan Kučera for the proof of Lem. 3.12, and for thorough and invaluable discussions about the Stickelberger ideal. Andrea Lesavourey is funded by the Direction Générale de l'Armement (Pôle de Recherche CYBER), with the support of Région Bretagne. This work is supported by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701) and by the PEPR quantique France 2030 programme (ANR-22-PETQ-0008).

References

Bab86. L. Babai: *On Lovász' lattice reduction and the nearest lattice point problem.* Combinatorica, **6**(1), pp. 1–13, 1986.

Bac90. É. Bach: *Explicit bounds for primality testing and related problems.* Math. Comp., **55**(191), pp. 355–380, 1990.

BCP97. W. Bosma, J. Cannon, C. Playoust: *The Magma algebra system. I. The user language.* J. Symbolic Comput., **24**(3-4), pp. 235–265, 1997, computational algebra and number theory (London, 1993).

BDF08. K. Belabas, F. Diaz y Diaz, E. Friedman: *Small generators of the ideal class group.* Math. Comput., **77**(262), pp. 1185–1197, 2008.

BDPW20. K. d. Boer, L. Ducas, A. Pellet-Mary, B. Wesolowski: *Random self-reducibility of Ideal-SVP via Arakelov random walks.* In *CRYPTO (2)*, vol. 12171 of *LNCS*, pp. 243–273, Springer, 2020.

BEF+17. J. Biasse, T. Espitau, P. Fouque, A. Gélin, P. Kirchner: *Computing generator in cyclotomic integer rings*. In *EUROCRYPT (1)*, vol. 10210 of *LNCS*, pp. 60–88, Springer, 2017.

BF14. J. Biasse, C. Fieker: *Subexponential class group and unit group computation in large degree number fields*. LMS J. Comp. Math., **17**(A), pp. 385–403, 2014.

BFHP21. J. Biasse, C. Fieker, T. Hofmann, A. Page: *Norm relations and computational problems in number fields*. arXiv:2002.12332v3 [math.NT], 2021.

BK21. O. Bernard, R. Kučera: *A short basis of the Stickelberger ideal of a cyclotomic field*. arXiv:2109.13329 [math.NT], 2021.

BL21. D. J. Bernstein, T. Lange: *Non-randomness of s-unit lattices*. Cryptology ePrint Archive, Report 2021/1428, 2021.

BLNR21. O. Bernard, A. Lesavourey, T. Nguyen, A. Roux-Langlois: *Log-$\mathcal{S}$-unit lattices using explicit Stickelberger generators to solve Approx Ideal-SVP (full version)*. Cryptology ePrint Archive, Report 2021/1384, 2021.

BPR04. J. Buhler, C. Pomerance, L. Robertson: *Heuristics for class numbers of prime-power real cyclotomic fields*. Fields Inst. Commun., **41**, pp. 149–157, 2004.

BR20. O. Bernard, A. Roux-Langlois: *Twisted-PHS: Using the product formula to solve Approx-SVP in ideal lattices*. In *ASIACRYPT*, vol. 12492 of *LNCS*, pp. 349–380, Springer, 2020.

BS16. J.-F. Biasse, F. Song: *Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields*. In *SODA*, pp. 893–902, SIAM, 2016.

CDPR16. R. Cramer, L. Ducas, C. Peikert, O. Regev: *Recovering short generators of principal ideals in cyclotomic rings*. In *EUROCRYPT (2)*, vol. 9666 of *LNCS*, pp. 559–585, Springer, 2016.

CDW17. R. Cramer, L. Ducas, B. Wesolowski: *Short Stickelberger class relations and application to Ideal-SVP*. In *EUROCRYPT (1)*, vol. 10210 of *LNCS*, pp. 324–348, Springer, 2017.

CDW21. R. Cramer, L. Ducas, B. Wesolowski: *Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time*. J. ACM, **68**(2), 2021.

CGS14. P. Campbell, M. Groves, D. Shepherd: *Soliloquy: A cautionary tale*, 2014, available at http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf.

Che13. Y. Chen: *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. Ph.D. thesis, Paris 7, 2013.

CN11. Y. Chen, P. Q. Nguyen: *BKZ 2.0: Better lattice security estimates*. In *ASIACRYPT*, vol. 7073 of *LNCS*, pp. 1–20, Springer, 2011.

DPW19. L. Ducas, M. Plançon, B. Wesolowski: *On the shortness of vectors to be found by the Ideal-SVP quantum algorithm*. In *CRYPTO (1)*, vol. 11692 of *LNCS*, pp. 322–351, Springer, 2019.

EHKS14. K. Eisenträger, S. Hallgren, A. Y. Kitaev, F. Song: *A quantum algorithm for computing the unit group of an arbitrary degree number field*. In *STOC*, pp. 293–302, ACM, 2014.

FpL16. FpLLL development team: *fplll, a lattice reduction library*, 2016, available at https://github.com/fplll/fplll.

GK89. R. Gold, J. Kim: *Bases for cyclotomic units*. Compos. Math., **71**(1), pp. 13–27, 1989.

GN08. N. Gama, P. Q. Nguyen: *Predicting lattice reduction*. In *EUROCRYPT*, vol. 4965 of *LNCS*, pp. 31–51, Springer, 2008.

HW38. G. H. Hardy, E. M. Wright: *An Introduction to the Theory of Numbers*. Oxford University Press, 1938, Fourth Edition.

Kuč86. R. Kučera: *On a certain subideal of the Stickelberger ideal of a cyclotomic field*. Archivum Mathematicum, **22**(1), pp. 7–19, 1986.

Kuč92. R. Kučera: *On bases of the Stickelberger ideal and of the group of circular units of a cyclotomic field*. J. Number Theory, **40**(3), pp. 284–316, 1992.

Laa16. T. Laarhoven: *Sieving for closest lattice vectors (with preprocessing)*. In *SAC*, vol. 10532 of *LNCS*, pp. 523–542, Springer, 2016.

LLL82. A. K. Lenstra, H. W. Lenstra, L. Lovász: *Factoring polynomials with rational coefficients*. Math. Ann., **261**, pp. 515–534, 1982.

LPR10. V. Lyubashevsky, C. Peikert, O. Regev: *On ideal lattices and learning with errors over rings*. In *EUROCRYPT*, vol. 6110 of *LNCS*, pp. 1–23, Springer, 2010.

MG02. D. Micciancio, S. Goldwasser: *Complexity of Lattice Problems*, vol. 671 of *The Kluwer International Series in Engineering and Computer Science*. Springer, 2002.

Mil14. J. C. Miller: *Class numbers of real cyclotomic fields of composite conductor*. LMS J. Comput. Math., **17**, pp. 404–417, 2014.

Nar04. W. Narkiewicz: *Elementary and Analytic Theory of Algebraic Numbers*. Springer Monographs in Mathematics, Springer, 3 edn., 2004.

Neu99. J. Neukirch: *Algebraic Number Theory*, vol. 322 of *Grundlehren des mathematischen Wissenschaften*. Springer, 1999.

PHS19. A. Pellet-Mary, G. Hanrot, D. Stehlé: *Approx-SVP in Ideal lattices with pre-processing*. In *EUROCRYPT (2)*, vol. 11477 of *LNCS*, pp. 685–716, Springer, 2019.

PZ89. M. Pohst, H. Zassenhaus: *Algorithmic Algebraic Number Theory*. Encyclop. Math. Appl., Cambridge University Press, 1989.

Sag20. Sage Developers: *SageMath, the Sage Mathematics Software System (Version 9.0)*, 2020, available at https://www.sagemath.org.

Sch87. C. Schnorr: *A hierarchy of polynomial time lattice basis reduction algorithms*. Theor. Comput. Sci., **53**, pp. 201–224, 1987.

Sch03. R. Schoof: *Class numbers of real cyclotomic fields of prime conductor*. Math. Comput., **72**(242), pp. 913–937, 2003.

Sch08. R. Schoof: *Catalan's Conjecture*. Universitext, Springer, 2008.

SE94. C. Schnorr, M. Euchner: *Lattice basis reduction: Improved practical algorithms and solving subset sum problems*. Math. Program., **66**, pp. 181–199, 1994.

Sin78. W. Sinnott: *On the Stickelberger ideal and the circular units of a cyclotomic field*. Ann. Math., **108**(1), pp. 107–134, 1978.

Sin80. W. Sinnott: *On the Stickelberger ideal and the circular units of an abelian field*. Invent. Math., **62**, pp. 181–234, 1980.

SSTX09. D. Stehlé, R. Steinfeld, K. Tanaka, K. Xagawa: *Efficient public key encryption based on ideal lattices*. In *ASIACRYPT*, vol. 5912 of *LNCS*, pp. 617–635, Springer, 2009.

Was97. L. C. Washington: *Introduction to Cyclotomic Fields*, vol. 83 of *Graduate Texts in Mathematics*. Springer, 2 edn., 1997.

Wes18. B. Wesolowski: *Generating subgroups of ray class groups with small prime ideals*. In *ANTS-XIII*, vol. 2 of *The Open Book Series*, pp. 461–478, Mathematical Sciences Publisher, 2018.